

AD A 026379

12

*Final Technical Report*

May 1976

## DEVELOPMENT OF AUTOMATED AIDS FOR DECISION ANALYSIS

By ALLEN C. MILLER, III  
MILEY W. MERKHOFFER  
RONALD A. HOWARD

JAMES E. MATHESON  
THOMAS R. RICE

*Prepared for*

THE DEFENSE ADVANCED RESEARCH PROJECTS AGENCY  
ARLINGTON, VIRGINIA

Contract Number MDA 903-74-C-0240

Amount of Contract \$155,000

DARPA Order Number 2742

Principal Investigator James E. Matheson  
(415) 326-6200, ext. 4450

Program Code Number PD420

Project Scientist Allen C. Miller, III  
(415) 326-6200, ext. 2865

Contractor Stanford Research Institute  
333 Ravenswood Avenue  
Menlo Park, California 94025

SRI Project Number 3309

Effective Date of Contract 1 June 1974

SRI Proposal Number MSU 73 133

Contract Expiration Date 30 September 1975

*Sponsored by*

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY  
DARPA ORDER NUMBER 2742

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U. S. Government.



**STANFORD RESEARCH INSTITUTE**  
Menlo Park, California 94025 • U.S.A.

Final Technical Report

May 1976

## DEVELOPMENT OF AUTOMATED AIDS FOR DECISION ANALYSIS.

By ALLEN C. MILLER, III  
MILEY W. MERKHOFF  
RONALD A. HOWARD

JAMES E. MATHESON  
THOMAS R. RICE

Prepared for

THE DEFENSE ADVANCED RESEARCH PROJECTS AGENCY  
ARLINGTON, VIRGINIA

Contract Number MDA 903-74-C-0240

SRI Project 3309

This research was supported by the Advanced Research Projects Agency of the Department of Defense under Contract No. MDA 903-74-C-0240

A

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER	2. GOV. ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle)  DEVELOPMENT OF AUTOMATED AIDS FOR DECISION ANALYSIS		5. TYPE OF REPORT & PERIOD COVERED  Final Technical Report	
7. AUTHOR(s) Allen C. Miller, III      James E. Matheson Miley W. Merkhofer      Thomas R. Rice Ronald A. Howard		6. PERFORMING ORG. REPORT NUMBER <del>MDA-903-74-C-0240</del> 3309	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Stanford Research Institute 333 Ravenswood Avenue Menlo Park, California 94025		8. CONTRACT OR GRANT NUMBER(s) MDA-903-74-C-0240 ✓	
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency Arlington, Virginia		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Defense Contract Administration Burlingame, California		12. REPORT DATE June 1976	13. NO. OF PAGES 224
		15. SECURITY CLASS. (of this report)  UNCLASSIFIED	
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this report)  Unlimited			
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from report)  Unlimited			
18. SUPPLEMENTARY NOTES			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)  Decision Analysis, Automated Decision Aids, Decision Models			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)  This report summarizes research undertaken to initiate development of a system of automated decision aids. The purpose of these aids is to facilitate the application of decision analysis to major decisions within the Department of Defense. The report contains: (1) a characterization of the spectrum of decision situations that are encountered in practice and an exploration of the implications of these characteristics for automated decision aids; (2) a description of the types of decision models available for analyzing various decision situations; (3) a description of the process by which decision models are constructed;			

**DD FORM 1473**  
1 JAN 73  
EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

CONTINUED

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

19. KEY WORDS (Continued)

20. ABSTRACT (Continued)

and (4) an identification of several modeling concepts that provide a basis for designing and constructing a pilot-level system of automated decision aids.

DD FORM 1473 (BACK)  
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)



## CONTENTS

LIST OF ILLUSTRATIONS . . . . .	vii
SUMMARY . . . . .	1
A. A Characterization of Decision Situations . . . . .	2
B. Types of Decision Models . . . . .	3
C. The Process of Modeling and Analyzing Decisions . . . . .	4
D. Concepts and Tools for Decision Modeling . . . . .	5
I INTRODUCTION . . . . .	9
A. The Purpose and Content of This Research . . . . .	9
B. The Importance and Difficulty of Creating a Decision Model . . . . .	11
C. The Scope of This Report . . . . .	12
D. Overview of the Report . . . . .	13
II TYPES OF DECISION PROBLEMS: IMPLICATIONS FOR MODELING . . . . .	15
A. Characteristics of Decisions and Decision Environments . . . . .	16
1. Time Dependence . . . . .	16
2. Sequence of Decisions and Information States . . . . .	17
3. Uncertainty . . . . .	18
4. Complexity . . . . .	19
5. Continuity . . . . .	20
6. Uniqueness . . . . .	21
7. Resources Affected by a Decision . . . . .	22
8. Scope of Decision . . . . .	22
9. Urgency . . . . .	23
B. Characteristics of Decision Makers . . . . .	24
1. Complexity of Preferences . . . . .	24
2. Resources Available for Analysis . . . . .	26
3. Experience and Training in Analyzing Decisions . . . . .	27
C. Characteristics of the Decision Process . . . . .	28
1. Number of People Involved in the Decision . . . . .	28

II	TYPES OF DECISION PROBLEMS: IMPLICATIONS FOR MODELING (Continued)	
2.	Organization and Interaction of People Involved in the Decision . . . . .	29
III	TYPES OF DECISION MODELS . . . . .	31
A.	Components of a Decision Model . . . . .	31
1.	System Variables . . . . .	32
2.	Sequence of Information States . . . . .	35
3.	Dependencies Among System Variables . . . . .	36
4.	Preference Model . . . . .	38
a.	Value Model . . . . .	40
b.	Risks Preference Model . . . . .	43
c.	Time Preference Model . . . . .	46
B.	Characteristics of Decision Models . . . . .	49
1.	Relative Emphasis on Explicit Modeling and Assessment . . . . .	49
2.	Model Complexity and Size . . . . .	50
3.	The Mathematical Structure of a Model . . . . .	51
IV	THE PROCESS OF MODELING AND ANALYZING DECISIONS . . . . .	55
A.	Modeling . . . . .	55
1.	Structuring . . . . .	57
2.	Assessment . . . . .	62
a.	Assessing State Variables . . . . .	63
b.	Assessing Dependencies . . . . .	65
c.	Assessing Preferences . . . . .	68
B.	Analysis . . . . .	74
1.	Model Solution and Decision Optimization . . . . .	75
2.	Sensitivity Analysis . . . . .	77
3.	Determining the Value of Information . . . . .	81
V	DECISION MODELING CONCEPTS . . . . .	85
A.	Basic Concepts: Entities and Operators . . . . .	86
1.	The Properties of Operators . . . . .	86
2.	Ambiguous Combinations of Entities of Operators . . . . .	89
3.	Connection Rules . . . . .	89
4.	Graphs and Trees . . . . .	90
B.	Function Graphs . . . . .	92

## V DECISION MODELING CONCEPTS (Continued)

1.	Using a Function Graph to Represent a Model for Choosing Between Alternative Missile Systems . . . .	94
2.	An Illustration of the Use of Function Graphs in a Computerized Decision Aiding System . . . . .	96
a.	Using CADMUS to Build a Model . . . . .	98
b.	The Completed Cost Model . . . . .	105
c.	Model Analysis . . . . .	110
C.	Decision Trees and Influence Diagrams . . . . .	115
1.	Properties of Decision Trees and Influence Diagrams . . . . .	118
a.	Probabilistic Independence . . . . .	118
b.	Expansion . . . . .	118
c.	Probability Trees . . . . .	119
d.	Decision Trees . . . . .	119
e.	Probability Assignment for Decision Trees . . .	121
f.	Influence Diagrams . . . . .	123
g.	Graphical Manipulation . . . . .	126
h.	Influence Diagrams with Decision Variables . .	129
i.	Node Terminology . . . . .	131
j.	Missing Influences . . . . .	132
k.	Relationship of Influence Diagrams to Decision Trees . . . . .	134
2.	An Example Use of the Influence Diagrams . . . . .	138
a.	Description of a Model Based on Influence Diagrams . . . . .	139
b.	Using the Model . . . . .	147

## APPENDICES

A	PROBABILISTIC DEPENDENCE . . . . .	153
1.	Independence Equation for Two Random Variables . . .	155
2.	Independence Equations for Three Random Variables .	156
3.	Relationships Among Independence Equations . . . .	158
4.	Combinations of Independence Equations . . . . .	160
5.	An Additional Relationship Among Independence Equations . . . . .	168
6.	Encoding the Twenty-Two Combinations of Independence Equations . . . . .	169
B	APPROXIMATE METHODS OF CALCULATING THE VALUE OF INFORMATION . . . . .	173
1.	Introduction . . . . .	175
2.	Computerizing the Single State Model . . . . .	177

## APPENDICES (Continued)

3.	Steps to Compute the Approximate Value of Information . . . . .	178
4.	Discrete Decision Variables . . . . .	179
5.	Continuous Decision Variables . . . . .	179
6.	Two-Stage Problems . . . . .	181
C	THE VALUE OF DECISION-DEPENDENT INFORMATION . . . . .	183
1.	The Primary Decision . . . . .	185
2.	Value of Perfect Information . . . . .	186
3.	The Concept of Clairvoyance . . . . .	190
4.	Practical Probability Assignment . . . . .	192
5.	Further implications of the Probability Assignment Model . . . . .	194
6.	Decision-Dependent Clairvoyance . . . . .	195
7.	Unequal Decision-Dependent Probabilities . . . . .	193
8.	Conclusion . . . . .	198
D	THE VALUE OF SEQUENTIAL INFORMATION . . . . .	201
E	THE VALUE OF FLEXIBILITY . . . . .	207
	REFERENCES . . . . .	213

## ILLUSTRATIONS

3.1	A Simple Decision Model . . . . .	39
3.2	A Decision Model With Structural and Preference Submodels . . . . .	39
3.3	A Decision Model with Value and Risk Preference Submodels . . . . .	40
3.4	A Decision Model with a Time Preference Submodel. . . . .	47
3.5	A Useful Type of Structural Model . . . . .	52
4.1	The Process of Modeling and Analyzing Decisions . . . . .	56
4.2	Indifference Curves on Surfaces . . . . .	71
4.3	Initial Model Progression in Decision Analysis . . . . .	79
4.4	Model Progression in Decision Analysis . . . . .	80
5.1	Illustration of Operators and Entities Using a Demand Curve . . . . .	87
5.2	Models of a Coal Burning Power Plant . . . . .	91
5.3	Computational Graphs . . . . .	93
5.4	Function Graph Representing Model for Choosing Between Proposed Missile Systems . . . . .	95
5.5	Control Panel used for Model Formulation by Cadmus . . . . .	99
5.6	First Step for Modeling the Missile Decision . . . . .	101
5.7	Second Step for Modeling the Missile Decision . . . . .	104
5.8	Third Step for Modeling the Missile Decision . . . . .	106
5.9	Fourth Step for Modeling the Missile Decision . . . . .	107
5.10	Final Step for Modeling the Missile Decision . . . . .	109
5.11	Sensitivity of Economic Cost to Number of Alerts per Year (Juggernaut System) . . . . .	112
5.12	Sensitivity of Economic Cost to Number of Alerts per Year (Ares System) . . . . .	113
5.13	Sensitivity of Economic Cost to Number of Alerts per Year (Juggernaut and Ares Systems) . . . . .	114
5.14	Probability Distribution Illustrating Uncertainty in the Number of Alerts per Year . . . . .	116

5.16	A Probability Tree . . . . .	120
5.17	A Generic Probability Tree . . . . .	120
5.18	A Generic Decision Tree . . . . .	121
5.19	A Four Node Decision Tree . . . . .	122
5.20	A Four Node Decision Tree Given the Assertion that $y$ will not Affect $x$ . . . . .	122
5.21	The Probability Assignment Tree . . . . .	122
5.22	Definitions Used in Influence Diagrams . . . . .	124
5.23	Two Node Influence Diagrams . . . . .	
5.24	Alternate Influence Diagrams for $x, y, a, S$ . . . . .	126
5.25	Graphical Manipulation of Influence Diagrams . . . . .	128
5.26	Graphical Manipulations Producing Non-Unique Results . . . . .	130
5.27	An Influence Diagram with Decision Nodes . . . . .	131
5.28	Some Sets Defined by the Node $g$ . . . . .	133
5.29	An Influence Diagram Representable by a Decision Tree . . . . .	135
5.30	Influence Diagram Requiring Probabilistic Manipulation Before Decision Tree Construction . . . . .	135
5.31	Influence Diagram Ready for Development into a Decision Tree . . . . .	136
5.32	The Process of Converting a Decision Network to a Decision Tree Network . . . . .	138
5.33	Development of Persian Gulf Scenarios . . . . .	145
5.34	Simplified Diagram of Persian Gulf Scenarios . . . . .	148
5.35	An Application of an Influence Diagram . . . . .	150
A.1	Combinations of Independence Equations . . . . .	161
B.1	Classification of Decision Problems . . . . .	176
B.2	Open and Closed Loop Sensitivities . . . . .	180
B.3	Compensation for a Discrete Decision Variable . . . . .	180
B.4	Computation for Continuous Decision Variables . . . . .	181
B.5	The Two Stage Quadratic Decision Problem . . . . .	182
C.1	The Primary Decision Tree . . . . .	187

C.2	Typical Decision Tree for Determining the Value of Perfect Information . . . . .	188
C.3	More Complex Tree for Determining the Value with Perfect Information . . . . .	189
C.4	Clairvoyant's Report Form . . . . .	190
C.5	Joint Distribution for Clairvoyant's Answers . . . . .	191
C.6	Decision Tree for Determining the Value of Complete Clairvoyance . . . . .	193
C.7	Probability Assignment Model . . . . .	194
C.8	Value of Clairvoyance on Launce Systems Only . . . . .	196
C.9	Value with Clairvoyance on Mars Only . . . . .	197
D.1	Sequential Information Decision Tree . . . . .	204
D.2	Decision Regions for Sequential Information . . . . .	206

## SUMMARY

The research summarized by this report was undertaken as a first step in the development of a system of automated decision aids designed to facilitate the application of decision analysis to major decisions within the Department of Defense. "Decision analysis", as used here, refers to a practical science of decision making that combines the fields of operations research and statistical decision theory. Applications of decision analysis to military problems have demonstrated its potential for assisting military decision makers. However future application will be limited unless automated aids are developed to speed up the decision analysis process and allow non-specialists access to the powerful problem solving tools that are currently available only to a few highly-trained and experienced decision analysts.

To advance the development of such a system of computerized decision aids, research was undertaken to develop a decision morphology: a precise characterization of the logical and analytic steps required to analyze a wide variety of decisions. This research has produced: (1) a characterization of the different kinds of decision situations that arise in practice and an exploration of the implications of these characteristics for automated decision aids; (2) a description of the types of decision models available for analyzing a variety of decision situations; (3) a description of the process of constructing decision models, and (4) an identification of several easily understood modeling concepts that provide a basis for designing and constructing a pilot-level system of automated decision aids. The paragraphs below provide brief summaries of the results in each of these four areas.



## A Characterization of Decision Situations

Decision situations differ in three basic respects, each of which affects the nature of appropriate modeling procedures and associated computerized decision aids. Decision problems can be characterized by (1) the nature of the decision environment, (2) the preferences and resources of decision makers, and (3) the process by which various individuals interact to reach a decision.

The applicability of various decision analytic procedures depends on the characteristics of the specific decision being analyzed. For instance, one important characteristic is the role time plays in the decision. If the basic elements of the decision environment or the relationships among these elements change significantly with time, an analyst must decide whether or not to use decision models with time-varying parameters. Since these models are much more difficult to construct and more expensive to analyze, dynamic models should be used only if the explicit representation of the time-varying aspects of the situations is critical to determining an optimal decision. Other relevant characteristics of a decision situation include: the degree of uncertainty present, which may dictate the use of probabilistic rather than deterministic models; the complexity and continuity of the environment, which influences the appropriate size of the model; whether the situation is unique or recurring; the resources affected; and the scope and urgency of the decision, which determine the appropriate level of resources that should be devoted to the development of a decision model.

The characteristics of decision makers define a second respect in which decision situations differ. For instance, decision makers can differ according to the complexity of their preferences. One commander might use a simple measure to guide his decision, such as whether or not his forces control a strategic location by nightfall. Another commander might base his decision on a variety of objectives, including not only whether his forces have gained control, but also the time required to do so, the casualties suffered, the psychological effect on the enemy, the morale of his troops, and so forth. Decision makers can also differ in their training

and experience with quantitative analysis. As a result, decision models must be adapted to produce results that can be easily understood and accepted by a particular decision maker. Similarly, the nature of decision models are governed by differences in the level of resources that decision makers can allocate to problem analysis.

The third major respect in which decision situations differ is the organizational structure within which decision makers must operate. Organizational structure can be characterized by the number of people involved in the decision and the nature of the interactions among them. When there is only one decision maker, a decision model will be concerned primarily with his alternatives, preferences, and information. However, when there are two or more decision makers, each must take into account the possible actions of the others. The manner in which interactions among decision makers are modeled strongly influences both the applicability and complexity of a decision model.

#### Types of Decision Models

For the purposes of this research, a decision model is any quantitative or logical abstraction of reality that is constructed and analyzed to help somebody reach a decision. The usual method of characterizing decision models is by the nature of their components. Virtually all decision models contain one or more of the following components: system variables, which represent elements of the decision environment; assumptions concerning the sequence or timing of new information about the decision environment; assumptions concerning the structural linkages or interdependences among system variables; and a specification of the decision maker's preferences among possible states of the world as represented by the system variables. The relative emphasis placed on these components determines the form of a particular model and the types of insights it can provide into a decision situation.

In addition, there are several global characteristics that distinguish decision models. One of these is the relative emphasis in the model on derivation rather than direct assessment of factors related to the decision.

In some situations it may be difficult to assess directly the implications of complex interactions among elements of the decision environment. However, doing so may significantly simplify an analysis. The size and complexity of a decision model are also important global characteristics that determine the model's applicability. Another common way of characterizing a decision model is in terms of its mathematical properties, such as whether the system variables are deterministic or probabilistic, continuous or discrete, static or time-dependent, and constrained or unconstrained. Although these characteristics need not be of great concern to the decision maker, they can be crucial in determining the feasibility of analyzing a model and its usefulness in describing the decision environment.

#### The Process of Modeling and Analyzing Decisions

To construct and analyze a decision model, it is necessary to carry out a number of logical procedures. They can be broken into two broad categories: modeling and analysis. The modeling process is primarily concerned with defining the various components of the decision model. The analysis process involves solving the model to determine its implications for decision making, and computing the effects of changing assumptions inherent in the model.

The modeling process is composed of two components: structuring and assessment. In structuring, the decision maker's alternatives, the possible outcomes that could result from his actions, the external factors affecting the outcome, and the linkages among the model variables are identified. During the assessment process the existing state of information about components of the decision model is quantified and the decision maker's preferences among possible outcomes, the time at which they might occur, and the risks associated with each alternative are specified.

The process of model analysis consists of model solution, sensitivity analysis, and determining the value of gathering additional information. The solution of the model consists of finding the optimal decision alternative and the corresponding probability distribution of possible outcomes. Sensitivity analysis is used to determine the relative importance of model

components and thus guides efforts to revise and improve the model. Determining the value of information provides the basis for decisions to gather new information, which can be used to update or restructure the model.

The various modeling and analysis procedures are used iteratively to develop a decision model. Thus at each stage in the development of a decision model, the model itself is used to determine which of its components should be modified and whether new information about the decision environment should be gathered.

### Concepts and Tools for Decision Modeling

The model building process often benefits from the use of graphical aids such as decision trees, flow charts, and block diagrams, and algebraic representations such as mathematical equations and computer programs. Therefore, an efficient decision aiding system should provide user communication in a language that facilitates the use of both graphical and algebraic aids. As a first step towards identifying such a language, research was conducted to determine the underlying concepts common to all decision model representations.

An important conclusion is that all models, regardless of whether they are represented in graphical or algebraic terms, are composed of two types of fundamental building blocks: "entities" and "operators". Entities are variables that describe the state of the decision environment or the decision maker's perception of the environment. Operators are directional functions that take entities as inputs and produce other entities as outputs. The purpose of operators is to describe the way in which entities are modified by the environment and the manner in which they are related. As an example, radar coverage and weather conditions might be entities in a decision model of an air strike. A functional relationship that relates weather conditions to radar coverage could be regarded as a model operator.

Decision models are produced by connecting appropriately defined entities and operators. In small models it is typical for specification of the links between operators and entities to be included in the definition

of the operators. However, the definitions of entities and operators, and the manner in which they are connected can be specified separately. Topological connections among operators and entities are specified by "connection rules". These rules determine exactly which entities are the inputs and outputs of each operator

When operators and entities are linked together, either individually or through the use of connection rules, the result is a computational graph. Computational graphs can be classified according to the degree to which they are interconnected. For example, the connections in some graphs contain loops. A computational graph contains a loop if two entities are connected by two or more paths with oppositely directed operators. If a computational graph contains no loops it is called a tree. An example is a decision tree. In a decision tree the entities are the nodes of the tree, the linkages are the branches, and the operators are the rules for evaluating the tree.

Specific types of graphs are useful as applied modeling tools. One such tool is a function graph--a graph composed of entities that are algebraic or logical variables, and operators expressed in the form of equations. Working with a decision maker, an analyst can construct a deterministic decision model in the form of a function graph. The advantage of working with a function graph is that the rules for its construction may be generated in an orderly way from answers the decision maker supplies to a set of specific and well-defined questions. A similar sort of graph, called an influence diagram, can be used to construct a model of the probabilistic relationships among elements of a decision problem. This modeling tool can be used to greatly reduce the number of probability assessments and amount of probabilistic processing required to analyze a decision model. Although further research is needed to gain a complete understanding of the properties of these applied modeling concepts, they have already proven useful in the analysis of several complex decision problems.

It appears that function graphs and influence diagrams may represent a significant advancement in the area of computer-aided decision modeling. These tools allow an individual to represent his or her knowledge regarding model structure in a way that is intuitive, and, at the

same time sufficiently precise to serve as a communication medium for a computer. Thus there is hope that computerized decision aiding systems can be developed that provide the military executive with powerful tools for improving his decision making ability without requiring him to learn specialized computer languages that are foreign to his present way of thinking.

## I INTRODUCTION

### A. The Purpose and Context of this Research

The research described in this report was conducted in response to a desire on the part of the Defense Advanced Research Projects Agency and the Decision Analysis Group of SRI to develop automated decision aids that would facilitate the application of decision analysis to major management decisions within the Department of Defense. When this research began, several computer programs had been written at SRI and elsewhere to deal with various aspects of analyzing decision problems, such as eliciting probability distributions or carrying out decision tree calculations. However, these computer programs dealt only with certain limited aspects of the decision analysis process and, for the most part, they required the user to be familiar with computer programming. The computer programs were not well suited to the needs of decision makers who lacked training in decision analysis and computer programming. In fact, computer programs developed by one group of decision analysts were seldom used by other analysts because the programs often contained procedures and limitations specific to the needs of the analysts who wrote them.

Since the SRI Decision Analysis Group had already written several general purpose computer programs for generating and analyzing decision trees (and other tree structures), it was necessary to decide whether the research would be devoted to expanding and generalizing the existing aids, or to investigating more generally the procedures used to describe and analyze decision problems and the ways in which automated decision aids could help analysts and decision makers carry out these procedures. It was decided that incremental improvements in the existing software for decision analysis were unlikely to produce decision aids beneficial to a wide variety of users and comprehensive enough to deal with all of the procedures used to analyze decision problems.

Therefore, this research project was undertaken, not with the aim of producing specific automated decision aids, but rather with the objective of characterizing all of the logical and analytical steps required to analyze a wide variety of decisions in a manner that would facilitate the design and implementation of automated decision aids. Since this characterization deals primarily with the general procedures used to analyze decisions, and not with specific computer programs, we have chosen to call it a "decision morphology" rather than use a more computer-oriented term like "algorithmic language".

To achieve the objectives of this research effort, a number of related research tasks were undertaken. The first task was a characterization of various possible decision environments and an exploration of the implications of different types of decisions for the design of appropriate automated decision aids. The next step was an investigation of the types of decision models available for analyzing the decision situations defined in the first task. The third task was an identification of the steps required to construct and analyze decision models. From the results of these three tasks, several general modeling concepts were identified which could provide the basis for a pilot-level system of automated decision aids.

Although the modeling concepts described in this report appear to provide a way to conceptualize and analyze a wide variety of decision problems, they should be applied to several realistic decision analyses to test their practicality and generality. With additional research it may be possible to generalize the specific modeling tools discussed in this report to the point where they will be appropriate for almost all decision modeling tasks. However, this may require the development of new tools for modeling and analysis to deal with decision situations for which the concepts described here are not adequate or appropriate. At the same time, future research should be directed toward unifying the modeling tools into a single, integrated system capable of dealing with all aspects of the decision modeling process.



## B. The Importance and Difficulty of Creating a Decision Model

Most of the existing computer programs for decision analysis deal with decision tree calculations, or the elicitation and processing of subjective probabilities. Since decision trees and probability elicitation are basic elements of decision analysis, it would seem logical that decision aids designed to help an analyst in these areas would be among the most useful and desirable. However, in applying decision analysis to many different types of decision problems, the SRI Decision Analysis Group has found that probabilistic processing and elicitation are not the areas in which an analyst or decision maker can benefit most from automated decision aids.

At the beginning of this research project each of the project leaders in the Decision Analysis Group was asked to review the projects that he had conducted and determine those aspects of the analysis that had been the most difficult and time-consuming. The consensus of this group was that the most difficult part of conducting a decision analysis--and the area that could benefit most from automated decision aids--was constructing a decision model that captured the essential elements of the decision under consideration. Probability elicitation, probabilistic processing, and the evaluation of decision trees all ranked below model development in terms of the amount of effort required to carry out the process and the potential usefulness of automated decision aids.

Admittedly, this informal survey covered only the experiences of one group of decision analysts. As a result, the importance assigned to decision modeling may be a reflection of the methodological approach of the SRI Decision Analysis Group as opposed to approaches used by other analysts. However, the survey covers a very broad range of decision problems, including decisions in both the public and private sectors, decisions characterized by varying degrees of complexity and uncertainty, and decisions ranging from specific resource allocations to broad questions of public policy. Furthermore recent research carried out by SRI for the Office of Naval Research on the development of operational decision aids for task force commanders has confirmed the importance of decision modeling in the analysis of military decision problems.

It is easy to see why the development of decision models plays an important part in the decision analysis process. Decision models allow analysts and decision makers to organize and rank in importance the many complex factors associated with major decisions. By considering each part of the model in turn an analyst can divide a major decision problem into a series of smaller, more manageable problems. At the same time a model can be used to determine the relative importance of the various elements of the decision problem. All too often decision makers become absorbed in one aspect of the decision problem and neglect other aspects that are equally, if not more, important. Even a simple decision model can show quickly which elements of the problem deserve the most attention by virtue of their relative influence on the final outcome. Perhaps equally important is the ability of a decision model to act as a vehicle for communication. An explicit model of all the elements of a complex decision allows everyone concerned to understand and contribute to the decision logic.

#### C. The Scope of this Report

This research report is not a textbook on decision analysis. Rather it is an exposition of those concepts, procedures, and models that should be included in a comprehensive system of decision aids. This report is concerned with the application of decision analysis rather than with the development of decision theory. Furthermore, it does not attempt to describe all the ways that people have conceptualized and applied decision analysis; rather it reflects the experience of the SRI Decision Analysis Group. However, the discussion of decision analysis contained in this report is intended to be sufficiently general to encompass most of the approaches taken to applying decision analysis.

It is assumed that the reader is familiar with the fundamentals of decision analysis. These are described in several publications [3,4,12,13]\*. However a detailed understanding of decision theory is not required to understand most of this material.

---

\*Figures in brackets correspond to the references.

#### D. Overview of the Research Report

Chapter II discusses the characteristics of different decision problems, decision makers, and decision processes, together with the implications of these characteristics for decision modeling. Chapter III outlines the different types of decision models, first by discussing the different components that can appear in various decision models and second in terms of more general characteristics of the entire model. Chapter IV deals with the process by which decision problems are modeled and analyzed. This is the process that will be facilitated by the development of automated decision aids. Chapter V contains the most novel findings of this research effort: several decision modeling concepts that can be understood by decision makers and are suitable for computer implementation. These conceptual tasks can be applied in the context of the decision modeling process described in Chapter IV to produce the types of decision models discussed in Chapter III. The concepts presented in Chapter V provide a basis for designing and constructing a pilot-level system of automated decision aids.

## II TYPES OF DECISION PROBLEMS: IMPLICATIONS FOR MODELING

Whether or not an automated decision aid is useful for analyzing a decision problem depends to a large extent on the characteristics of the decision being analyzed and the decision maker for whom the analysis is conducted. It is doubtful that any single decision aid will be equally applicable to all of the decision problems that one might wish to analyze. Furthermore, a decision aid that is helpful to one decision maker might not be useful to a decision maker who has a different level of experience in analyzing decisions or who must operate within a different organizational structure. Developing a comprehensive system of automated decision aids requires understanding the range of decision problems to which it would be applied.

It would be difficult to enumerate all of the different types of decision problems that one might wish to analyze. However, there are certain characteristics of decision problems and decision makers that clearly influence the appropriateness of various modeling procedures and related computerized decision aids. These characteristics and their implications for modeling and analyzing decisions are discussed in this section.

In attempting to define the different types of decision problems, one must be careful to distinguish between the characteristics of the decision environment and the characteristics of the model that is developed to analyze the decision. A statement such as, "That is a linear programming problem" may be more descriptive of the model used to analyze a decision than it is of the decision environment. Since models are by definition abstractions of reality, it is quite appropriate that a decision model and the situation it represents have different properties. For example, an inherently non-linear decision problem could be represented by a linear model. The discussion that follows deals with the characteristics of decision environments, not decision models, although some of the same characteristics could be used to describe models.

The list of characteristics described in this section is not exhaustive. There are many ways to characterize decision problems and decision makers, but some attributes are more useful descriptors than others. Taken as a whole, the following characteristics provide the dimensions for describing any particular decision problem.

A. Characteristics of Decisions and Decision Environments

Some of the major determinants of the appropriateness of various decision analytic procedures are the physical characteristics of the environment as they are perceived by the decision maker. A decision model--and the procedure used to build it--need not have exactly the same characteristics as the decision environment, but the two obviously should be closely related. The following sections outline the implications of various characteristics of the decision environment for developing an appropriate model.

1. Time Dependence (Static vs. Dynamic Environments):

There are many decision problems in which the basic elements of the decision environment, and even the relationships among these elements, change with time. Models with time-varying parameters can be used to represent such situations, but dynamic models are inherently more difficult to analyze than static models. When attempting to model a time-varying decision environment, an analyst must determine whether variations with time must be represented explicitly or whether the time-dependent characteristics of the environment can be averaged or aggregated so that they can be represented by a static model. Fortunately, dynamic decision problems can often be adequately approximated by models with static parameters or functional relationships. One way to simplify the time dependence in a model is to represent time as a series of discrete intervals. In that way the value of model parameters in each time interval can be considered as separate quantities related to each other by functions that describe how they change from one time period to another.

In some ways time can be considered a model parameter just like any other element of the model. The dependence of some quantity on time can be modeled in the same manner as its dependence on other model parameters

Assumptions about the level of detail required to specify time are similar to those for other model parameters.

However, time possesses certain characteristics that distinguish it from other model parameters. It provides a natural index with which other model parameters can be identified. It also defines the order in which decisions are made and states of information are changed when these elements are functions of time.

2. Sequence of Decisions and Information States (Sequential vs. Nonsequential Decision Environments):

Even if the major elements of a decision problem and their interrelationships are not functions of time, it may be necessary to investigate the implications of making a decision based on several possible states of information. For example, a military commander may decide on the best way to position his forces without knowledge of the enemy's position, and then find it necessary to reconsider the same decision after receiving new information from his intelligence sources. In fact, he could find himself making a whole sequence of resource allocation and reallocation decisions based on changing states of information, even though the configuration of enemy forces was not actually changing with time. Thus, changing states of information can occur in both static and dynamic decision environments.

Changes in the state of information can affect both the estimates of parameters in a decision model and the structure of the model itself. Models in which the structural relationships among parameters depend on changing states of information tend to be complicated and difficult to specify.

Explicit consideration of changing information states is usually required in a decision model when there is a sequence of decisions to be made at the same time that the information is changing. In this case it is necessary to analyze each decision for each of the possible information states that might exist when the decision must be made. Sequential decisions can be modeled efficiently if one can identify repetitive or similar information states and combine them to reduce the size of the model.

Sequences of decisions can themselves lead to changing states of information, because the decisions produce new information about the environment or cause significant changes in the environment itself. Decisions that produce new information about the environment are generally information-purchasing decisions, such as initiating an intelligence activity or conducting research to determine the operating characteristics of a new weapons system. Other decisions are made with the intent of directly influencing the environment and producing specific results. As changes in the decision environment take place, the decision maker often receives new information that changes his view of the problem.

3. Uncertainty (Probabilistic vs. Deterministic Decision Environments):

Decision makers almost never know the exact consequences of choosing an alternative at the time that they make a decision. In an attempt to reduce the uncertainty associated with their decisions, decision makers will often allocate significant portions of their resources to gathering information.

Many attempts have been made to analyze uncertain decision environments with deterministic models. However, the use of deterministic models often leaves the decision maker wondering whether they adequately describe all of the contingencies that might occur. Deterministic models are particularly inappropriate for information-purchasing decisions, since the objective of these decisions is to reduce uncertainty about the environment. Deterministic models do a poor job of measuring changes in the level of uncertainty.

Since probabilistic models are more difficult to process than deterministic models, the question that must be answered in any decision modeling effort is: To what extent should uncertainty be modeled explicitly? Although various modeling techniques have been developed to help answer this question (such as assessing probability distributions only for those model variables that have the greatest effect on the model's output), the appropriate level of probabilistic modeling depends

on the time and resources available for modeling as well as on the nature of the decision environment. Often a deterministic model can be used to gain an initial understanding of a decision problem before a more complete probabilistic model is developed. In fact, deterministic models often evolve into probabilistic models of the same situation as model variables and relationships are redefined in probabilistic terms.

One of the primary reasons for modeling uncertainty explicitly is that human intuition is notoriously poor at determining the effect of interactions among uncertain quantities. Individuals tend to think about decision environments in deterministic terms, even though they recognize that there are important elements of the problem that are beyond their control and about which they have limited information. For this reason the probabilistic model can produce insights that are not obvious from intuitive reasoning.

#### 4. Complexity (Simple vs. Complex Decision Environments):

Decision environments can range from those with a few basic elements to those comprised of large numbers of interrelated components. However, complex decision problems do not always lead to complicated decision models. Often many of the components inherent in a decision situation can be aggregated to form the variables of a simple model. For example, economists regularly model the many millions of business transactions that occur each day in the national economy in terms of a few aggregate parameters like gross national product, personal consumption, and investment.

Trying to determine the proper level of simplification and aggregation to use in a decision model is probably the most difficult part of modeling. On the one hand, simplified models are relatively easy to analyze and explain to others. However, the more specific elements of the decision environment are aggregated into a few global parameters, the more difficult it becomes to define exactly what those parameters mean. Furthermore, oversimplifying a decision model can destroy one of the model's primary benefits: its ability to break a decision



down into its component parts and let the decision maker deal separately with each of the components.

The complexity inherent in a decision situation can itself be a source of uncertainty. If it is necessary to conceptualize and model the situation in terms of a relatively small number of aggregate parameters, the decision maker may find himself uncertain about the manner in which the aggregate parameters interact even though he might be quite knowledgeable about the many detailed interactions that occur in the environment. In the process of developing a simplified model of a complex decision situation it may be necessary to introduce uncertainty into the model to avoid a more detailed analysis.

5. Continuity (Continuous vs. Discrete Decision Environments):

In some decision situations there are only a few discrete alternatives available. For example, an officer in charge of air defense might have to decide which of three available weapons systems to assign to an attacking enemy aircraft. In other situations, the decision maker can select any value over a continuous range, for example when one is determining the proper elevation of an artillery piece. Similarly, elements of the decision environment that are beyond the decision maker's control can be classified according to whether they take on only a few discrete values or can vary over a specified range.

In reality, the distinction between continuous and discrete quantities is somewhat blurred. When trying to decide what percentage of his forces to hold in reserve, the commander of a large force might have so many discrete alternatives that his choice is essentially one of selecting a fraction of his force on a continuous scale of zero to 100%. This distinction may seem subtle, but it has a major impact on the nature of the model used to analyze the decision.

If the decision environment is characterized by essentially continuous quantities, such as economic costs and benefits, the relationships among these quantities can often be modeled by continuous functions or equations. Alternatively, if major elements of the decision

environment are inherently discrete, their interrelationships are usually represented by discrete operators such as matrices, difference equations, or a detailed enumeration of the values of a certain parameter given all of the possible combinations of the variables upon which it depends. A decision tree can be viewed as an example of the latter type of transformation.

Whether a particular decision is modeled with continuous or discrete quantities depends in part of the level of aggregation inherent in the model. One way to simplify a decision model is to approximate the continuous quantity with a variable that takes on only a few representative values. For example, it might be sufficient to model possible states of the weather in terms of sunshine, rain, or snow even though there are many possible amounts and types of precipitation. Decision environments characterized by a high degree of complexity are often represented by discrete decision models in order to minimize the size of the model and the analytical effort required to solve it.

6. Uniqueness (Unique vs. Repetitive Decision Environments):

The applicability of various decision modeling techniques is partially determined by the frequency with which similar decisions must be made. Many of the methods of classical statistics may be invoked when dealing with repetitive decisions, such as when to reorder parts for an inventory. Furthermore, the common elements of repetitive decisions can be analyzed and modeled prior to the time that a decision needs to be made. This prior analysis makes it possible to specify and even automate decision rules so that the decision can be made on a routine basis. As a result, repetitive decision models tend to be delegated to subordinates even though the decisions may be extremely important to the organization. The allocation of weapons systems to fleet defense is an example of an important, repetitive decision that is delegated to specially-trained, subordinate decision makers.

Unique decisions require both the development of a new decision model and a representation of states of information based on little,

if any, statistical data. In this case both the model structure and the values of model variables must rely to a large extent on subjective estimates. The analysis of unique decisions is primarily focused on gaining an understanding of the problem rather than fine-tuning a model to the point where it can be used routinely and efficiently for dealing with similar decision situations.

7. Resources Affected by a Decision (Major vs. Minor Decisions):

The extent to which a decision model is developed is partially determined by the scale of resources affected by the decision. There are two resource levels that characterize the decision environment: the range of outcomes that might result from the decision, and the level of resources allocated by the decision maker. The level of resources allocated need not determine the importance of the outcomes. Even if the decision maker has limited resources at his disposal, he may be in a position to influence events of major importance. For example, a small but strategically located military unit may have more influence on the outcome of a battle than other units with more resources at their disposal in inferior positions.

It is worthwhile expanding a decision model whenever the cost of the additional effort is outweighed by the benefits that are likely to occur from a deeper understanding of the decision environment. As the scale of resources affected by the decision goes up, so do the benefits that can be expected to accrue from the insights produced by additional modeling. Thus, formal decision models generally find application in the analysis of major decisions involving significant levels of resources, while decisions with less important consequences tend to be made on an intuitive basis.

8. Scope of Decision (Specific vs. Policy Decisions):

Some decisions deal with specific resource allocations. Other decisions are designed to set policy and give guidance to those who must make specific decisions. Models of policy decisions tend to be at a more aggregate and less detailed level than those of the specific decisions

for which they provide policy guidelines. However, in the course of modeling a policy decision, it may be necessary to analyze specific decisions that might be affected. For example, when setting policy as to how his forces should respond to various enemy actions, a commander may wish to analyze briefly the decisions that may be faced by his subordinates.

The effect of policy decisions is to define the objectives that should be used by others in analyzing specific decisions, thereby simplifying the decision models needed for the relevant specific decisions. As a result, models of policy decisions must often incorporate a more detailed consideration of complex preference issues than do models of specific decisions.

Policy decisions are often used to coordinate the decision-making activities of several individuals or organizational subdivisions. This means that models of policy decisions often must take into account the interactions of several decision makers. This requirement tends to complicate the analysis and expand the size of the decision model, even though the model deals primarily with aggregate parameters. Furthermore, since policy decisions are usually designed to give guidance to a number of subordinate decision makers, it is usually necessary to model explicitly the sequence of information states associated with specific decisions. This requirement also tends to expand the size of models dealing with policy decisions.

9. Urgency (Planning vs. Crisis Decision Environments):

The time available for making a decision has a significant effect on the kind of modeling that is possible. If a decision must be made immediately, detailed explicit modeling is not possible. However, the question of how soon a decision must be made is itself usually a decision. Even though the cost of delaying a decision is high, it may be wise to spend some time carefully considering its implications. A bad decision made hastily may be worse than a delayed decision.

It is often possible to anticipate and even model at least some elements of a crisis decision before the crisis occurs. Many crisis

decisions can be reduced to manageable proportions if those elements of the problem that can be anticipated are subjected to careful analysis before the crisis occurs. For example, a naval commander considering the possibility of an enemy submarine attack could develop a model of the decision he would face in that situation, including such known parameters as the speed and armaments of his ships. The use of such a partially pre-structured model would make it possible for him to analyze quickly a specific combat situation by updating the model with a description of the situation at hand.

There will always be some urgent decisions that cannot be anticipated. Because of the limited time available, models for such decisions tend to be composed of a small number of aggregate parameters that describe only the most important elements of the problem. The main purpose of such models is not to give the decision maker a detailed understanding of the consequences of the decision, but rather to help him sort out the major elements of the decision and be sure that nothing important is overlooked.

#### B. Characteristics of Decision Makers

Since decision models are relevant only if they are put in terms that are meaningful to the decision maker, the characteristics of the decision maker are as important as those of the decision environment in determining the nature of an appropriate model. The implications of various characteristics of decision makers for modeling and analyzing decisions are discussed in the following paragraphs.

##### 1. Complexity of Preferences (Simple vs. Complex Preferences):

A problem that often characterizes decisions, especially those in the public sector, is the problem of determining the criteria with which to judge the outcome of a decision. Analyses of financial decisions are often based on the assumption that the objective is to minimize the net cost or maximize the net profit that results from a particular decision. Yet quite often the financial decision maker is concerned about many other things besides profit and loss: the effect of the decision on

the organization's public image, the internal politics of the organization, the decision's effect on morale, and the possibility that a bad outcome could jeopardize the decision maker's career. Obviously preferences like these are difficult to quantify. Military and public sector decision problems, which often involve such basic considerations as the value of human life, contain preferences that are even more difficult to quantify.

However, the fact that preferences are difficult to quantify does not mean that it is impossible to model them. Fortunately, preference modeling has been the object of considerable research over the past several decades. The result of this research has been to identify various types of preferences (preferences among various attributes of possible decision outcomes, time preference, and risk preference) and to develop procedures for assessing and combining each of these preference types within a single decision model. In decision problems containing relatively simple and easily identifiable objectives, it is often unnecessary to model explicitly certain types of preferences (such as time and risk preference). The subject of preference modeling is discussed in more detail in Chapters III and IV of this report.

In many public sector decision problems, the difficulty associated with modeling preferences is compounded by the lack of a clearly identified decision maker or decision-making procedure. Decisions made "in committee" or as a result of competition among various agencies may require a rather detailed assessment and modeling of the preferences of all of the competing parties in order to identify areas of possible compromise and focus the debate on the underlying differences and objectives. A decision analysis need not explicitly model the adversary process to help in its resolution. Usually a decision model can be constructed based on the assumption of a single set of preferences, and then each of the interested parties can exercise the model using his own preference structure. Typically, this sort of analysis demonstrates that certain preference differences are not worth arguing about because they lead to the same decision. (Explicit modeling of the process by which several decision makers resolve their differences is another level of analysis which is discussed separately below in Section C, Characteristics of the Decision Process.)

## 2. Resources Available for Analysis (Limited vs. Extensive Analytical Capability):

Closely related to the scale of resources affected by a decision is the level of resources made available for analyzing the decision. The difference between the two quantities is that the level of resources affected by the decision is usually determined by the decision environment while the resources available for analysis can be determined by the decision maker. Determining the level of resources to allocate to the analysis of a decision is itself a secondary decision which could be analyzed. A tertiary decision would then be the determination of how much time and effort to devote to the secondary decision, and so on. In most situations it is doubtful that a detailed analysis of these second- and third-order decisions is warranted. As a rule of thumb an individual might set a policy that for every major decision he will use some percentage (say one percent) of the resources to be allocated by that decision to determine the best method for making the allocation.

There are several types of resources that a decision maker could devote to the analysis of a decision problem. To a certain extent each type of resource is needed for the analysis, and the level of each resource allocated has a significant effect on the nature of the resulting decision model. The types of resources include: the decision maker's personal time, the time of individuals with expertise in various aspects of the particular decision problem at hand, the assistance of analysts with training in modeling decisions, the involvement of competing or cooperating decision makers, and the support of data-processing facilities.

The application of these resources to the analysis of a decision problem allows the decision maker to consider, and possibly model, the elements of a decision problem in greater detail. However caution should be exercised in developing detailed models and analyses of decision problems. All too often, detailed models are constructed to analyze elements of a decision problem that do not have a significant effect on the decision outcome. As resources are allocated for modeling a decision situation in detail, it is important that these resources be constantly focused on the most important elements of the decision problem.

3. Experience and Training in Analyzing Decisions (Intuitive vs. Analytical Decision Making):

The process by which a decision is reached depends on the backgrounds of the individuals involved in the decision. Some people prefer to reason in verbal, non-quantitative terms; to them, a detailed quantitative model might be inappropriate even if the resources to construct one were available. Other decision makers find quantitative models and analyses useful as a check for their own intuitive reasoning. Still others find them the only satisfying decision procedure. An important requirement for constructing a model to assist a decision maker with little experience in analyzing decisions is that the conclusions derived from the model be expressed in terms that are meaningful to the decision maker. This usually means that model parameters must be carefully defined to match the way that the decision maker thinks about the problem and that complicated interactions within the decision model must be summarized in intuitive terms.

The types of decision models appropriate to a particular situation depend not only on the analytical training of the decision maker, but also on the training of the individuals supplying information to the decision maker. If the people with expertise about the decision environment have not been trained in the procedures for estimating subjective probability distributions, the development of a probabilistic model requiring such distributions as inputs would necessitate the training of the various subject experts before the model could be used.

When it comes to modeling, a little training may be worse than none at all. All too often decision models are developed that depend more on an individual understanding of certain analytical techniques than on the nature of the decision problem under investigation. A decision maker may unconsciously exclude certain types of models from consideration, even though they might be appropriate for describing a particular decision environment, because he was either unaware of their existence or intent on structuring the model in such a way that a particular analytical technique would be applicable. As decision makers and analysts gain experience in modeling a wide variety of decision problems, they tend to develop models



that rely less on specific procedures or techniques and more on the nature of the problem being analyzed.

### C. Characteristics of the Decision Process

The term "decision process" refers to the interaction of the individuals associated with making a decision. Even if there is only one decision maker in a particular situation, his decision may be influenced strongly by the information supplied by other individuals or by the social or organizational structure within which he finds himself. Obviously the process by which people interact in order to reach decisions has a strong influence on the types of decision models appropriate for analyzing decisions.

#### 1. Number of People Involved in the Decision (Individual vs. Multi-person Decision Making):

Although it is common to refer to the "decision maker," major decisions are seldom made by a single individual in isolation. Even though one person may appear to be the decision maker, his options and objectives may be strongly influenced by the decisions of others, especially when his decisions must conform with organizational policy. Furthermore, the information upon which the decision is based and the manner in which it is analyzed may be the result of interaction between a decision maker and a large number of other individuals. The manner in which the decision is modeled will depend on the number of decision makers, experts (suppliers of information), and analysts involved in the decision-making process.

When there is only one decision maker, a decision model will be concerned primarily with his alternatives, preferences and information about the environment; in this case, the environment is not trying to anticipate his actions. However, when there are two decision makers, each must take into account the possible actions of the other. Each decision maker is faced with the task of estimating not only his own state of knowledge but also that of the other individual, who in turn must assess the amount that the first decision maker knows about his

information, and so on. Attempting to model situations where "I know that he knows that I know that he knows, etc." is a very difficult process unless a number of simplifying assumptions are built into the model. One assumption, which leads to the standard game-theory approach to decision making, is that the decision makers share the same information about the environment and are capable of analyzing all of the possible interactions that could occur between them. More complicated and realistic game-theoretic approaches to multi-person decision making assume that the decision makers have different states of information which are unknown to their competitors, but these approaches often produce open-ended models that are impossible to analyze.

When there are more than two decision makers who can affect the same outcome, modeling the decision process becomes even more difficult because there are now enough people involved to permit the formation of coalitions. In addition to making his primary decision, each individual must make a set of secondary decisions concerned with reaching agreements with other groups of decision makers. Although conceptually straightforward, models of coalitions among decision makers tend to be enormously complicated and difficult to solve.

Fortunately decision situations in which there is only one decision maker with several people supplying information about the decision environment are considerably easier to model than those with multiple decision makers. When there are several, possibly conflicting sources of information about the environment, a decision model should be designed to allow each piece of information to be incorporated into the analysis. While this may produce a somewhat more complicated model, the resolution of conflicting pieces of information does not require major changes in the types of models used to analyze decisions.

## 2. Organization and Interaction of People Involved in the Decision (Structured vs. Non-structured Decision Processes):

The analysis of situations involving several decision makers is greatly simplified if the decision makers find themselves in an organizational or economic environment that governs or limits the manner in

which they interact with each other. For example, in an economic market composed of many decision makers, none of whom controls a sufficiently large portion of the market to determine the prices of goods and services, the decisions of one individual are effectively isolated from those of another, although collectively their decisions determine the price and quantity of every item in the market.

There are a number of different ways in which the interactions among decision makers are controlled. In hierarchical organizations, it is safe to assume that an individual's decision will override those of others lower in the organization. While this assumption may not always hold in practice, it is usually sufficient for the development of decision models. Economic interactions are somewhat less structured but relatively easy to define. The degree to which one decision maker must take into account the decisions of others depends on such factors as their relative shares of the market and the relative costs of various outcomes for each decision maker. The interactions of political decisions are often the least structured and the most difficult to make. Models of political decisions usually contain limiting assumptions about the extent to which each decision maker tries to outguess the others, and as a result these models tend to be very crude approximation of the political process.

### III TYPES OF DECISION MODELS

This section deals with the various types of models that can be used to analyze decisions. The characteristics of different decision models are discussed both in terms of the components contained in the model and its general mathematical structure. The purpose of this section is not to describe specific models but rather to present an overview of the general characteristics common to most decision models.

For the purposes of this discussion, a decision model is any quantitative or logical abstraction of reality that is constructed and analyzed to help somebody reach a decision. This definition is broad enough to include almost any mathematical representation of reality, and appropriately so, for almost every effort to develop a quantitative model is ultimately justified by its potential for aiding some form of decision making. Even models that do not contain an explicit representation of the decision-making process are often used to gain insight into the environment surrounding a decision.

#### A. Components of a Decision Model

Decision models can be characterized by the types of components they contain, by the relative importance of various components, and by the way in which components are represented and interpreted.

The types of components in any decision model can be classified as follows: system variables, which represent elements of the decision environment; the sequence or timing of information about the system variables; the structural linkages or interdependencies among the system variables; and a specification of the decision maker's preferences among possible states of the world as represented by the system variables. Each type of component may have different names in different types of models, but this classification is sufficiently general to encompass the elements of any decision model. The various components are described in more detail in the following sections.

## 1. System Variables

System variables are those quantities chosen to represent elements of a decision environment that can exist in several possible states. Identifying and defining all of the system variables contained in a model determines the scope and specificity of the model: the range of possible situations that the model can represent and the degree of detail with which the model approximates reality. Specifying a particular value for each system variable completely defines the state of the environment as represented by the model. The only variables in the model that are not system variables are those that describe the decision maker's preferences -- preferences among states of the environment specified by the values of the system variables.

System variables can be divided into two categories: those that are directly and completely controlled by the decision maker, and those that are not. System variables that are completely controlled by the decision maker are called decision variables. All other system variables, including those that are influenced by the decision maker's actions but are not completely under his control, are called state variables (or environmental variables) since they define the state of the decision environment.

Decision variables must be defined in such a way that they represent the alternatives available to the decision maker, rather than outcomes that are only partially under his control. For example, a businessman may or may not be able to control the price at which he sells his product depending on the nature of the economic market within which he operates. A monopolist is free to specify the price of his products; for him, price is a decision variable. A businessman in a market characterized by many producers (such as a farmer selling grain) may have to accept the price determined by the marketplace; for him price is a state variable. In some cases, a system variable can either be a decision variable or a state variable depending on how it is used in the model. For example, a farmer has the option of demanding any price he wishes for his grain, knowing that if the price he specifies exceeds the market price he will not find a buyer. A model that included the farmer's decision whether

to set the price of his grain at the prevailing market price would contain the farmer's price as a decision variable.

State variables define all of the elements of a decision model other than a decision maker's alternatives and preferences. State variables can be determined in a two-step process. First, the variable itself is identified and defined. The definition includes the range of possible values that the variable can assume and the scale along which these values are measured. Second, the value of a state variable is specified, either at a particular level or in terms of a set of possible values with associated probabilities. The value of a state variable need not be known at the time that a model is constructed. In fact, one of the purposes of the model may be to determine the value of certain state variables. However, a major portion of the modeling process may be devoted to defining state variables carefully.

State variables can be classified as aleatory or fixed depending on whether the decision maker's state of information about a variable is represented by a probability distribution or by a single deterministic value. The decision maker's state of information and the relative importance of a state variable in determining the output of the model govern whether a state variable is represented probabilistically or deterministically. Since it is relatively difficult to process aleatory state variables, decision models typically include the assumption that all but the most important state variables can be represented by a fixed or deterministic value.

State variables can also be classified as either exogenous or endogenous. Exogenous state variables are the inputs to the model. They are not influenced by the values of other system variables, and they must be specified before the model can be analyzed. Endogenous state variables are those whose values or probability distributions are determined by the model; they are usually not specified by the user. In some cases an endogenous state variable will be specified by the user in order to exercise a subsection of the model for which the variable is an input. It is also possible for an internal or endogenous state variable to become exogenous by reversing the linkages or transformations among state

variables. The reversal of transformations within a model is discussed in more detail in the following sections.

A third way of classifying state variables is by whether or not they are affected by the decisions described in the model. State variables that are not affected by the actions of the decision maker are included in the model to specify possible states of information or situations in which the decision maker will have to select a course of action. State variables that are affected by a decision maker's choices are used in a model to represent the way in which a decision maker interacts with his environment. Since state variables that are affected by the actions of the decision maker are a function of decision variables, they must be endogenous or internal state variables.

A fourth way of classifying state variables is according to whether or not their values are known with certainty when any of the decisions represented in the model are made. A revealed state variable is one whose value is known prior to the selection of a course of action specified by one of the model's decision variables. Unrevealed state variables are specified only in terms of probability distributions at each of the decision points in the model. Unrevealed state variables are included in decision models primarily to specify the range of outcomes that could occur as a result of the actions taken by a decision maker. Specifying the things that might happen after a course of action has been selected facilitates the assessment of a decision maker's preferences by allowing him to separate his state of information about the occurrence of various outcomes from his preferences among the outcomes.

Those system variables that are direct inputs to a model of the decision maker's preferences are called outcome variables. Outcome variables can be either state or decision variables. The possible values that these variables can assume are called outcomes. Outcomes are those quantities that a decision maker would like to know in order to determine the desirability of a particular state of the world as represented in the model. Outcome variables are typically determined by asking the decision maker a question such as, "If you had to leave on a long vacation shortly after making a decision and could not find out what had happened until many years

later, which quantities would you most like to know in order to see how things had turned out?"

## 2. Sequence of Information States

One of the essential elements of a probabilistic decision model is the order in which information is revealed about aleatory state variables. Since the order in which information is revealed controls a decision maker's state of information, it also influences his decisions. Specifying the sequence of information states in a decision model defines the order in which aleatory state variables are revealed. This specification does not define the effect that revealing a state variable has on the probabilistic description of other state variables contained in the model. The specification of the sequence of information states in a decision model can be made independently of an assessment of the manner in which new information changes the decision maker's state of information.

Defining the sequence of information states in a decision model is not the same as specifying the time dependence of the system variables and their relationships to each other. For example, if the initial position and velocity of a spacecraft is known, its trajectory can be calculated as a function of time. If the decision maker is uncertain about the spacecraft's initial position and velocity, his state of information about its position at any future time can be represented by a probability distribution. A model of the decision to make corrections in the course of the spacecraft might include specification of the times at which the decision maker will receive new information about the actual position of the spacecraft. The timing of this new information may be completely independent of the time dependence of the spacecraft's trajectory, and even independent of the decision maker's earlier states of information about the trajectory. Thus the sequence of variable revelation and their time dependence can be specified independently.

The order in which aleatory state variables are revealed can affect the decision process if their revelation is interspersed with a series of decisions. In this case a change in the sequence of information states can affect the decision maker's best course of action. On



the other hand, if the order in which two aleatory state variables are revealed is reversed, but no decisions are made between the time that the two variables are revealed, the change in the sequence of information states will not lead to a new decision.

The order in which information is revealed in a decision model need not be the same as the order in which information is assessed. For example, we might assess the likelihood that a forecaster will correctly predict a rainy day even though a decision is based on the likelihood of rain given the weather forecast. The rules of probability theory can be used to transform states of information assessed in one order into those that occur if the sequence of events is altered. As a result, decision models may require the specification of two sequences of information states, one for assessment and one for use in the model.

### 3. Dependencies Among System Variables

A decision model must contain information about the manner in which system variables are linked together. Like system variables, the structural linkages can be either static or dynamic and either deterministic or probabilistic. Furthermore, the relationships among system variables, like the variables themselves, can be defined in two steps: an identification of the fact that one set of system variables depends on another, and a specification of the way in which the values of one set of system variables can be derived from the other.

For example, a decision model could contain a relationship giving the position of an aircraft in terms of its location at takeoff and its velocity and heading after takeoff. The exact transformation by which the location of the aircraft is derived from the other quantities need not be specified to show the relationship among the system variables. In fact, the user might want to experiment with defining the transformation at various levels of complexity and realism. For instance the transformation giving the position of an aircraft in terms of its initial position, velocity and heading may or may not be designed to take into account the curvature of the earth.

If all of the system variables related by a particular transformation are deterministic, then the transformation itself must be deterministic. A deterministic transformation specifies the exact value of the dependent state variables for each combination of possible values for the independent state variables. Deterministic transformations can also be used to define the structural linkages among probabilistic or aleatory state variables. In this case the rules of probability theory are used in conjunction with the deterministic transformation to produce the probability distribution for the dependent state variables from the probability distribution of the independent state variables.

However, uncertainty in the outputs of a transformation may be caused not by uncertainty in the inputs to the transformation, but rather by uncertainty in the transformation itself. Probabilistic dependencies among state variables are defined in terms of conditional probability distributions. However, conditional probability distributions are used sparingly in decision models because they require a great deal of effort to specify, analyze and interpret. The nature of probabilistic dependence is a very complicated subject, especially when there are several dependent variables. This subject is explored briefly in Appendix A.

The distinction between state variables and the transformations that link them together becomes somewhat clouded when the transformation is parameterized. For example, one of the structural linkages in a decision model might define one state variable as a multiple of another, without specifying the exact value of the multiplier. In this case, the multiplier may be varied, or even defined probabilistically, in order to see how changes in its value affect the output of the model. If the model is formulated in this manner, the multiplier should be considered a state variable rather than a part of the transformation. The transformation should be redefined to produce the value of the dependent state variable by taking the product of the multiplier and the other independent state variable.

The distinction between state variables and transformations is harder to recognize when the transformations are probabilistic because

both state variables and model linkages may be defined in terms of probability distributions. The difference between them is that transformations are defined by conditional probability distributions (distributions for the output given the value of the input) while the probability distributions for state variables are not conditional.

#### 4. Preference Model

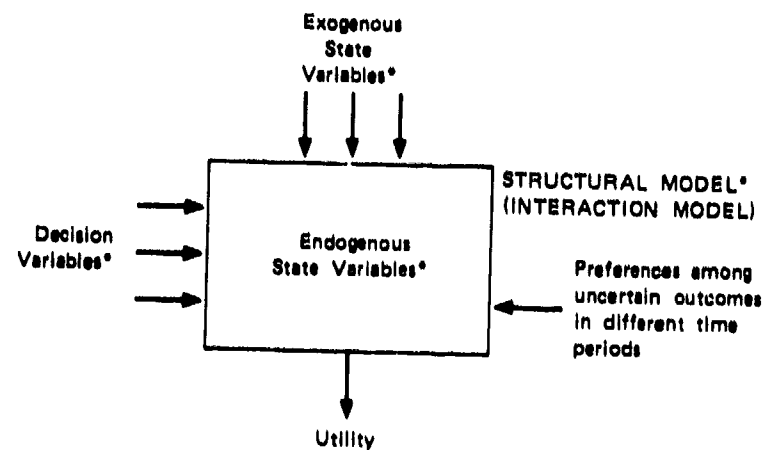
A decision model can be described as a mapping\* from decisions and possible states of the decision environment to a measure of preference that indicates how well the decision maker likes the resulting outcomes. This measure of preference is called utility. This view of a decision model is shown in Figure 3.1. However, it is useful to think of a decision model as being composed of two submodels: a structural model that describes the interactions among system variables, and an overall preference model that describes the decision maker's preferences among the outcomes determined by the structural model. This decomposition of a decision model into a structural model linked to a preference model is shown in Figure 3.2.

Outcome variables, those state variables on which the decision maker defines his preferences, can be probabilistic or deterministic and static or time-dependent. Thus an overall preference model may have to include the decision maker's attitude toward uncertain or risky situations and his preferences for outcomes occurring at different points of time.

In simple decision models the two submodels may be indistinguishable. For instance, decisions are sometimes modeled with interaction matrices which show the decision maker's alternatives, possible states of nature (combinations of state variable values), and the utility associated with each decision/state variable pair. In this case, outcome variables are not represented explicitly.

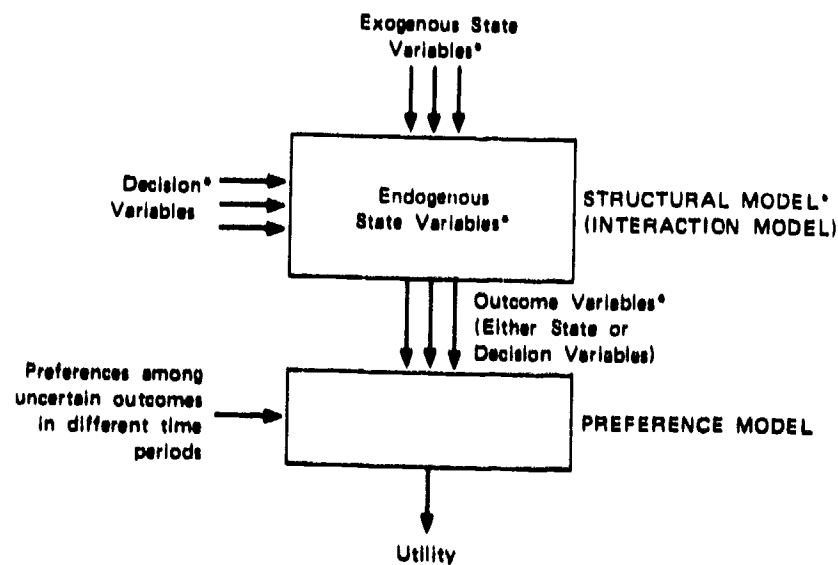
---

\*The phrase "mapping from A to B" refers to the mathematical process of transforming elements in A into elements in B through some functional relationship.



\*May be a function of time.

FIGURE 3.1 A SIMPLE DECISION MODEL



\*May be a function of time.

FIGURE 3.2 A DECISION MODEL WITH STRUCTURAL AND PREFERENCE SUBMODELS

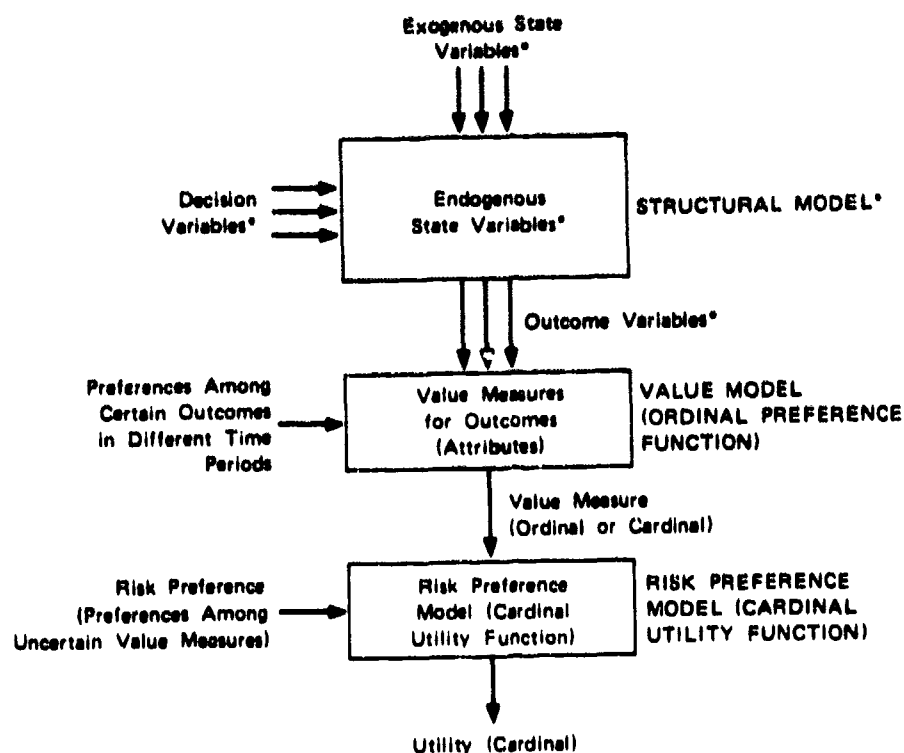
However, most decision models are composed of separate structural and preference submodels. Because of the importance and difficulty associated with modeling a decision maker's preferences, an overall preference model is often separated into a set of submodels. There are several ways to accomplish this decomposition based on certain assumptions about the form of the decision maker's preferences. The three main types of submodels used to specify preferences--value models, risk preference models, and time preference models--are described below.

a. Value Model

A value model specifies a decision maker's preferences among deterministic outcomes that can occur in several time periods. A value model does not specify the decision maker's attitude toward uncertainty or risk; it must be combined with a risk preference model to form a complete model of the decision maker's preferences. However, in many decision environments, the range of possible outcomes is such that the decision maker is willing to evaluate uncertain situations in terms of their expected values. In this case, specification of the value model is all that is needed to define the decision maker's preferences.

The advantage of separating an overall preference model into a value model and a risk preference model, as shown in Figure 3.3, is that it simplifies the process of specifying the decision maker's preferences. If his preferences are determined without the use of a distinct value model, the decision maker must state his choices not only among combinations of possible outcomes, but also among all possible probability distributions over the outcomes.

The process of specifying choices among deterministic combinations of outcomes is difficult enough without including uncertainty. Without a separate value model, the decision maker must answer questions like, "Would you prefer a situation in which you had an equal chance of destroying either 30% or 70% of the enemy force while losing 20% of your own force, or a situation in which you destroyed 50% of the enemy force and had an equal chance of losing 10% or 30% of your own force?" By



\*May be a function of time.

FIGURE 3.3 A DECISION MODEL WITH VALUE AND RISK PREFERENCE SUBMODELS

using a separate value model the decision maker can answer somewhat simpler questions, such as, "Would you prefer a situation in which you had destroyed 50% of the enemy force and lost 20% of your own forces, or a situation in which you had destroyed 30% of the enemy's forces and lost only 10% of your own forces"? While the second question is not easy to answer, it is certainly easier to think about than the first question.

A value model transforms deterministic outcomes, which may or may not be functions of time, into a one-dimensional, static value measure. A value model includes not only the decision maker's preferences among different possible outcomes, but also his preferences among outcomes occurring at different times. Thus it includes the decision maker's attitude toward changing the occurrence of an outcome from one time period to another. As long as certain restrictive assumptions are built into

the preference model, a value model can be constructed so that a decision maker's time preference is modeled separately.

Depending on how a value model is defined, the value measure that it produces may be either an ordinal or cardinal quantity. If a decision maker is asked to specify only his preferences among possible combinations of outcomes, or to identify sets of outcomes among which he is indifferent, then the value measure is an ordinal quantity. In other words, the fact that one combination of outcomes produces a higher value measure than another means only that the first is preferred to the second. The relative levels of the two value measures do not indicate how much more one set of outcomes is worth than another.

On the other hand, if a decision maker specifies his preferences among combinations of outcomes by specifying how much more one set of outcomes is worth than another (measured in terms of some well-defined unit, like money) then the value measure is a cardinal quantity. This means that the value measure can be used to determine, not only which combinations of outcomes are preferred, but also the relative values of various combinations of outcomes, measured in the units of the value measure. For financial and economic decisions, it is often possible to specify a cardinal value measure, but for decision problems with many diverse outcomes it is usually necessary that the value measure be an ordinal quantity.

The simplest way to construct a value model is by simply enumerating all possible combinations of outcomes and then specifying the value measure associated with each combination. Often this process is simplified by determining the decision maker's preferences among a few possible sets of outcomes and then using this information to draw approximate indifference curves in the space spanned by the outcome variables. Once the general form of the indifference curves has been ascertained, they can be approximated by equations that map the outcomes into the value measure.

One procedure for constructing a value model, called the "multi-attribute" approach, is to measure separately the decision maker's preferences for various levels of each outcome variable or attribute. This

Thus the multi-attribute approach need not restrict the form of the value model except for the requirement that the value measure be a cardinal quantity.

However, the multi-attribute approach is usually accompanied by the assumption that the mapping,  $X[\cdot]$ , from the individual value measure to the composite value measure is linear. With this assumption the value model is restricted to transformations of the following form:

$$V = a_1 W[O_1(t)] + a_2 W[O_2(t)]$$

There are many possible value models that cannot be represented in this form. For example, this assumption would preclude a value model that specified that the value measure was equal to the product of the various outcomes at a specific point in time.

#### b. Risk Preference Model

When a value model is used to specify the decision maker's preferences among deterministic outcomes, a separate risk preference model must be developed to represent the decision maker's attitude toward uncertainty and risk. As shown in Figure 3.3, a value model transforms the outcomes of the structural model into a one-dimensional, static value measure. However, since the outcomes are generally uncertain, so is the value measure. A risk preference model transforms the probability distribution over the value measure into a single deterministic number, called utility, that measures the decision maker's attitude toward uncertainty in the value measure.

Utility is a cardinal quantity that can be processed algebraically, but only in certain ways. The properties of utility are defined by a set of axioms\* that govern the manner in which it can be processed mathematically. The implication of these axioms is that, if a decision maker prefers one distribution of uncertain value measures to another, the preferred distribution is assigned a higher utility. Furthermore,

---

\*These axioms are described in detail in [5], [8], and [11].



produces an individual value measure for each outcome variable. The individual value measures associated with the outcome variable are then combined to produce a single, composite value measure. Since the value measure associated with each outcome variable is combined algebraically with the other value measures, it must be defined as a cardinal quantity. This means that the relative value of two possible levels of an outcome variable must be defined in terms of some measurable quantity like dollars, years, or lives. Unfortunately some attempts to construct a value model using a multi-attribute approach have used procedures that produce ordinal value measures associated with each of the outcomes; these individual value measures are then combined algebraically as if they were cardinal quantities. The composite value measure produced by a multi-attribute value model has meaning only if the individual value measures are cardinal quantities.

In general, the multi-attribute approach does not restrict the form of a value model, except that the value measure produced by the model and the value measures associated with each of the outcome variables must be cardinal quantities. Regardless of how it is developed, a value model can be viewed as a mapping  $V[.]$ , from the outcomes,  $O(t) = [O_1(t), O_2(t), \dots]$ , to the value measure,  $v$ . This mapping can be written  $v = V[O(t)]$ . With the multi-attribute approach each of the outcomes,  $O_1(t)$ , is mapped onto a corresponding value measure,  $w_1$ . Each of these mappings can be written:  $w_1 = W_1[O_1(t)]$ . The composite value measure is derived from the individual value measures with another transformation:

$$v = X[w_1, w_2, \dots] = [w_1[O_1(t)], w_2[O_2(t)], \dots]$$

As long as the transformations,  $w_1[.]$ , from the outcomes to the individual value measures are all one-to-one mappings, it is always possible to define a transformation,  $X[.]$ , from the individual value measures to the composite value measure such that:

$$V[O_1(t), O_2(t), \dots] = X[W_1[O_1(t)], W_2[O_2(t)], \dots]$$

the utility of any distribution of values is equal to the expected utility\*. Although it is permissible to calculate the expected value of the various utilities that might result from an uncertain situation, the utility axioms do not allow other mathematical manipulations, such as taking the difference of the expected utilities associated with two uncertain situations. Since the utility axioms are equally applicable to preferences among uncertain combinations of outcomes and preferences in terms of an uncertain value measure, the mathematical restrictions on how utilities can be processed apply whether or not a preference model is used as a whole or decomposed into a value model and a risk preference model.

A risk preference model, as shown in Figure 3.3, is a mapping from a value measure to utility. The nature of this transformation, especially the nonlinearities inherent in the transformation, define the decision maker's attitude toward uncertainty and risk. This transformation can be written:  $u = R[v]$ , where  $v$  is the value measure,  $u$  is the utility and  $R[.]$  is the mapping defined by a risk preference model.

A risk preference model is combined with a value model that maps the outcomes into the value measure:  $v = V[O(t)]$ . Together, these two submodels form a preference model which defines a mapping,  $U[.]$ , from the outcomes to utility:  $u = U[O(t)]$ . The use of distinct value and risk preference models does not restrict the form of an overall preference model as long as  $U[.]$  and  $V[.]$  produce the same ordinal ranking among the possible sets of outcomes. In other words, if an overall preference model,  $U[.]$ , and a value model,  $V[.]$ , represent the same preferences when the outcomes are known with certainty, then it is always possible to define a risk preference model,  $R[.]$ , such that the combination of the value model and the risk preference model,  $R[V[.]]$  produce the same preferences for uncertain outcomes as the overall preference model,  $U[.]$ .

---

\*The expected value of an uncertain quantity is equal to the sum of all possible values, each multiplied by its associated probability. Thus expected utility is the sum of all possible utilities, each multiplied by its associated probability.

There only are two assumptions inherent in constructing an overall preference model from a value model and a risk preference model: (1) the decision maker's attitude toward risk is defined over all of the possible outcomes that could occur, including outcomes that occur in different time periods, and (2) it does not change with time. These assumptions are necessary because information about the time at which various outcomes occur is not contained in the value measure, the input to a risk preference model.

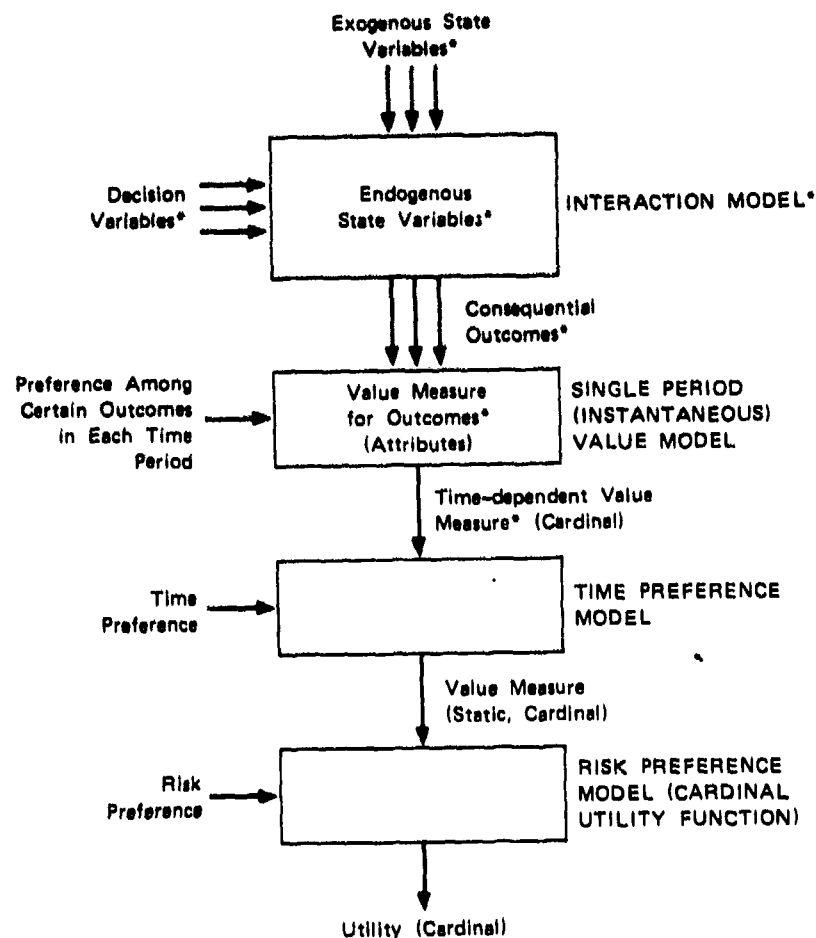
c. Time Preference Model

A preference model can be expanded even further by constructing a separate submodel of the decision maker's time preferences. This is accomplished by decomposing the value model shown in Figure 3.3 into two submodels: a single-period (instantaneous) value model\* and a time preference model. These sub-models are shown in Figure 3.4. The single period value model maps the outcome at each point in time into a corresponding level of a time-dependent value measure. This value measure, unlike the one discussed previously can have a different value at each point in time. The time preference model maps the time-dependent value measure into a static value measure that includes the decision maker's preferences for the occurrence of outcomes in different time periods. The single-period value model produces the same mapping from the outcomes to the time-dependent value measures at each point in time since all of the decision maker's preferences with respect to time are included in the time preference model.

Although the static value measure can be either an ordinal or cardinal quantity, the time-dependent value measure must be a cardinal quantity. If the time-dependent value measure were an ordinal quantity that ranked the decision maker's preference in any time period, but not his preferences among outcomes in different time periods, it would be impossible for the time preference model to establish preferences among

---

\*Time is treated as a discrete quantity in this section, but the concepts apply equally well to models where time is a continuous quantity.



\* May be a function of time.

FIGURE 3.4 A DECISION MODEL WITH A TIME PREFERENCE SUBMODEL

the outcomes in different time periods. For example, given that a decision maker prefers outcome A to outcome B in either of two time periods and that he prefers either A or B to occur in the first time period rather than the second, it is not possible to establish whether he prefers outcome B in period 1 to outcome A in period 2 unless further information is supplied. Thus the single-period value model must measure all of the possible combinations of outcomes at a particular point in time along a cardinal value scale so that these value measures can be manipulated by the time preference model to produce a static value measure.

In Figure 3.3 the value model maps the time-dependent outcomes into a static value measure:  $v = V[O(t)]$ . In Figure 3.4 this mapping is accomplished in two steps. The single-period value model,  $S[.]$ , maps the outcomes at each point of time into the time-dependent value measure,  $s(t)$ . This mapping can be written:  $s(t) = S[O(t)]$ . The time preference model,  $T[.]$ , maps the time-dependent value measure into the static value measure:  $v = T[s(t)]$ . Thus the mapping performed by the overall value model,  $V[.]$ , can be written:

$$V[.] = T[S[.]].$$

In general, decomposition of the value model into a single-period value model and time preference model need not restrict the form of the value model, except that both the time-dependent and static value measures must be cardinal quantities. However, the assumption that the time preference model has a particular form, such as the sum of the time-dependent value measures for each time period weighted by an appropriate discount factor, does restrict the form of the value model.

A specification of the decision maker's time preference, either in a time preference model or as an integral part of an overall preference model, is sometimes used in place of an explicit model of the decision maker's resource allocation decisions over time. In this case the decision maker's preferences may be based on both the time at which an outcome occurs and the time at which information about the outcome is revealed. For example, an individual faced with an equal chance of winning or losing a large sum of money one year from now might prefer to resolve the uncertainty now even though the time at which the gain or loss occurs cannot be changed. If he knows in advance whether he will win or lose the money, he can modify his current resource allocation decisions accordingly. If the resource allocation decisions are not an explicit part of a structural model, but are modeled implicitly as part of the decision maker's time preference, the sequence and timing of his states of information must be defined as an outcome. Otherwise the preference model will be based on insufficient information.

## B. Characteristics of Decision Models

The characteristics of a decision model are only partially defined by the nature of the model's components. In addition, decision models can be described by their overall structure and scope. The purpose of this section is to describe the global characteristics of various types of decision models.

### 1. Relative Emphasis on Explicit Modeling and Assessment

By definition, models are incomplete representations of reality. Therefore, whenever a decision maker's preferences or state of information are assessed, additional models could be developed to help determine the assessed quantities. The designer of a decision model will always be faced with a choice between direct assessment and further modeling. As a result, decision models can be characterized by the relative emphasis placed on assessment and modeling.

Simple models usually rely heavily on direct assessments of the decision maker's preferences and state of information about the environment. For example, the interaction matrix discussed in the preceding section is a relatively simple model to construct, but it requires that the decision maker state his preferences for all possible combinations of decisions and states of the environment. This type of model requires the decision maker to aggregate mentally the effects of the interactions among his decisions and elements of the environment at the time he assesses his preferences. If the decision environment is complex, the decision maker may be very uncomfortable about making such global assessments. Even if the decision maker feels that he can make the assessments required by a simple decision model, he should proceed with caution. It is easy to demonstrate that humans are very poor at assessing the interactions of probabilistic quantities.

As a model is developed, preliminary assessments may be used temporarily to represent undeveloped subsections of the model. This procedure allows an analysis of the simplified model to determine the importance of each subsection of the model and the relative advantage

of developing a detailed model of that subsection. This process also makes it possible to determine the appropriate level of assessment to use in place of further modeling.

## 2. Model Complexity and Size

Another way of characterizing decision models is by the number of system variables contained in the model and the intricacy of the interactions among them. The complexity of a model is related to, but not determined by, its size. Even a small decision model, one with relatively few system variables, can contain a very complicated set of interdependencies among variables. For instance, situations in which the decision maker's state of information about the environment changes rapidly and unpredictably often require a complex decision model even though the environment can be adequately described by a few aggregate variables.

The level of aggregation used to define system variables will affect both the size and complexity of a decision model. Model designers are often faced with a choice between using a large number of specific but sparsely-interconnected system variables or a smaller number of aggregate variables that are highly interdependent because they represent all of the relationships among the more specific variables. The appropriate level of aggregation for system variables is often determined by the difficulty associated with assessing the interdependencies among them. If system variables are overly aggregated, the decision maker may find it impossible to assess intuitively the intricate interactions among them. In this case it is necessary to disaggregate some of the system variables to simplify the task of assessing the interdependencies among them.

Decision makers have a great deal of difficulty comprehending and accepting results derived from a large, complex decision model if the interaction of system variables cannot be explained in a simple, logical manner. This often leads the model designer to make assumptions about the degree to which system variables are interdependent which may or may not reflect reality. Often the most limiting assumptions in a decision model are those related to the interactions among system variables,

rather than assumptions about the definitions or values of variables. In modeling efforts where the interactions among system variables have been oversimplified, it may be useful to expand the size of the model by defining a larger number of more specific variables whose interdependencies are easier to define.

### 3. The Mathematical Structure of a Model

One of the most common ways of characterizing a decision model is in terms of its mathematical properties, such as whether the system variables are deterministic or probabilistic, continuous or discrete, static or time-dependent, and constrained or unconstrained. While these characteristics may not be of great concern to a decision maker, they can be crucial in determining the feasibility of analyzing the model and its usefulness in describing the decision environment.

Deterministic models are those in which it is assumed that the state variables and the relationships among them are known with certainty. Probabilistic models contain uncertain state variables, and may or may not contain uncertain structural linkages among the system variables. The simplest probabilistic models are those with few, if any, probabilistic relationships among the system variables. Models with a large number of probabilistic dependencies are generally very difficult to specify and analyze.

System variables and structural linkages can also be either continuous or discrete. Models can be defined with discrete system variables and continuous transformations, but not with continuous system variables and discrete transformations. It is often desirable to treat time as a discrete quantity and specify probability distributions in terms of a discrete set of possible outcomes in order to facilitate the analysis of a decision model with a digital computer.

Two types of decision models that have proven especially useful in the work of the SRI Decision Analysis Group are: (1) those with continuous state variables, some of which are probabilistic, and a set of relationships among the system variables represented primarily by continuous, deterministic transformations; and (2) those with discrete, probabilistic state variables and probabilistic linkages among the state variables.



In the first type of decision model, the structural model is divided into two submodels as shown in Figure 3.5. Probabilistic dependencies among the state variables, if any exist, are represented by an

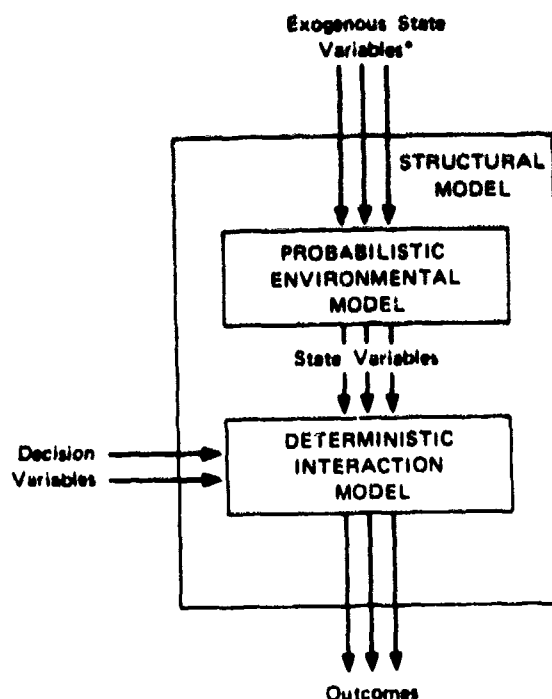


FIGURE 3.5 A USEFUL TYPE OF STRUCTURAL MODEL

environmental model which is not affected by the actions of the decision maker. The state variables produced by the environmental model are combined with the decision variables in a deterministic interaction model which produces the outcomes. One of the assumptions contained in this type of model is that the actions of the decision maker do not influence his state of knowledge about the state variables. This assumption, plus the decomposition of the structural model into separate probabilistic and deterministic submodels, makes this class of models relatively easy to define and analyze.

The second type of decision model, containing probabilistic relationships among discrete, probabilistic system variables, is usually represented by decision trees. The fact that the variables are discrete means that their possible values can be represented by branches on a decision tree, and the probabilistic relationships among variables are easy to specify in terms of probabilities associated with various branches of a tree.

Decision models can also be characterized by whether or not system variables and the relationships among them are functions of time. It is common for state variables to be time-dependent, but it is quite rare for the transformations among them to vary with time. This is partially due to the fact that a model can often be constructed so that the time-varying portions of transformations are treated as state variables. This is done in the same way that the uncertainty inherent in a transformation is represented by a state variable. For example, when one of the dependencies in a model specifies that one state variable is a time-varying multiple of another, the multiplier can be considered a time-dependent state variable and the transformation can be defined as one which produces the dependent state variable from the multiplier and the independent state variable.

Time, like any other model parameter, can be either continuous or discrete. If time is modeled as a discrete variable and only a few time periods are considered, it is possible to specify the occurrence of model elements in different time periods as separate system variables. In this case the time dependency of the model is suppressed and each of the system variables is a stationary quantity. This allows the user to model a time-varying decision environment with a decision model that is essentially static.

Another mathematical property that characterizes decision models is whether or not the values of the system variables are constrained to lie in specified intervals. The inclusion of constraints has a significant effect on the form of a decision model, in many cases simplifying the task of analysis. Constraints are an essential feature of certain techniques of mathematical analysis, such as linear programming. However, constraints are often crude approximations of reality, and they may exclude from consideration courses of action that are both feasible and desirable. For example, a model of a decision to allocate aircraft maintenance tasks to various repair facilities might be based on the assumption that each facility has a certain maximum capacity. The solution of such a model might overlook the fact that by slightly

expanding the capacity of one of the repair facilities the work could be carried out in a more cost-effective and efficient manner.

In place of constraints, decision models can incorporate a description of the cost involved in allowing system variables to take on values beyond their normal operating ranges. In the example discussed above, the model could include the cost of expanding the capacities of the various repair facilities. Obviously this would require a somewhat more complicated decision model, but the model would more accurately reflect the true nature of the decision.

The appropriateness of using constraints in decision models is influenced by the organizational structure within which the decision maker must operate. Policy decisions in large organizations are often stated in the form of constraints, even though subordinate decision makers might take courses of action more in keeping with the goals of the organization if policy were specified in terms of tradeoffs rather than constraints. Policy constraints supply guidance to a decision maker when he is operating near one of the constraint boundaries, but otherwise they are not helpful. Policy constraints are particularly inappropriate when they are applied to state variables instead of decision variables. For instance, issuing an order to a military commander to avoid contact with the enemy while maintaining a particular position overlooks the fact that the outcome of contact with the enemy is partially determined by factors beyond the commander's control.

Often the mathematical structure of a decision model is described by the particular solution or optimization technique used to analyze the model. For example, one might refer to a particular decision analysis as a linear programming model, a Lagrange multiplier model, or a decision tree model. While such characterizations are often useful for describing the model's properties, it is unfortunately the case that many models are constructed specifically to make use of a particular solution technique. The process of modeling, which will be discussed in detail in the following section, should be oriented primarily toward developing an adequate and efficient representation of a decision situation, and only secondarily to producing a mathematical structure amenable to analysis by a particular mathematical technique.

#### IV THE PROCESS OF MODELING AND ANALYZING DECISIONS

To construct and analyze a decision model, it is necessary to carry out a number of logical procedures. These procedures are discussed in this chapter. They can be broken into two broad categories: modeling and analysis. The modeling process is primarily concerned with defining the various components of a decision model described in the preceding section. The analysis process involves solving a model to determine its implications for decision making, and computing the effects of changing assumptions inherent in the model.

Figure 4.1 shows an overview of the process of modeling and analyzing decisions. This figure illustrates the fact that the elements of modeling (structuring and assessment) and the elements of analysis (model solution, sensitivity analysis, and determining the value of information) are used iteratively to develop a decision model. In other words, the model itself is used to determine which of its components should be modified, expanded, or contracted. Furthermore, the model is used to guide efforts to gather new information which, in turn, may lead to a modification of the structural definitions or assessed values contained in the model.

Each of the steps shown in Figure 4.1 is discussed separately in this section. However, none of them can be carried out independently: each requires information and insight produced by the others.

##### A. Modeling

The elements of a decision situation can be modeled in two steps: identification of model components and specification of the decision maker's state of information about them. Thus the process of modeling decisions can be divided into two steps: structuring and assessment. Structuring consists of identifying the elements of a decision model, and assessment consists of quantitatively specifying the values of variables and the nature of transformations among them.

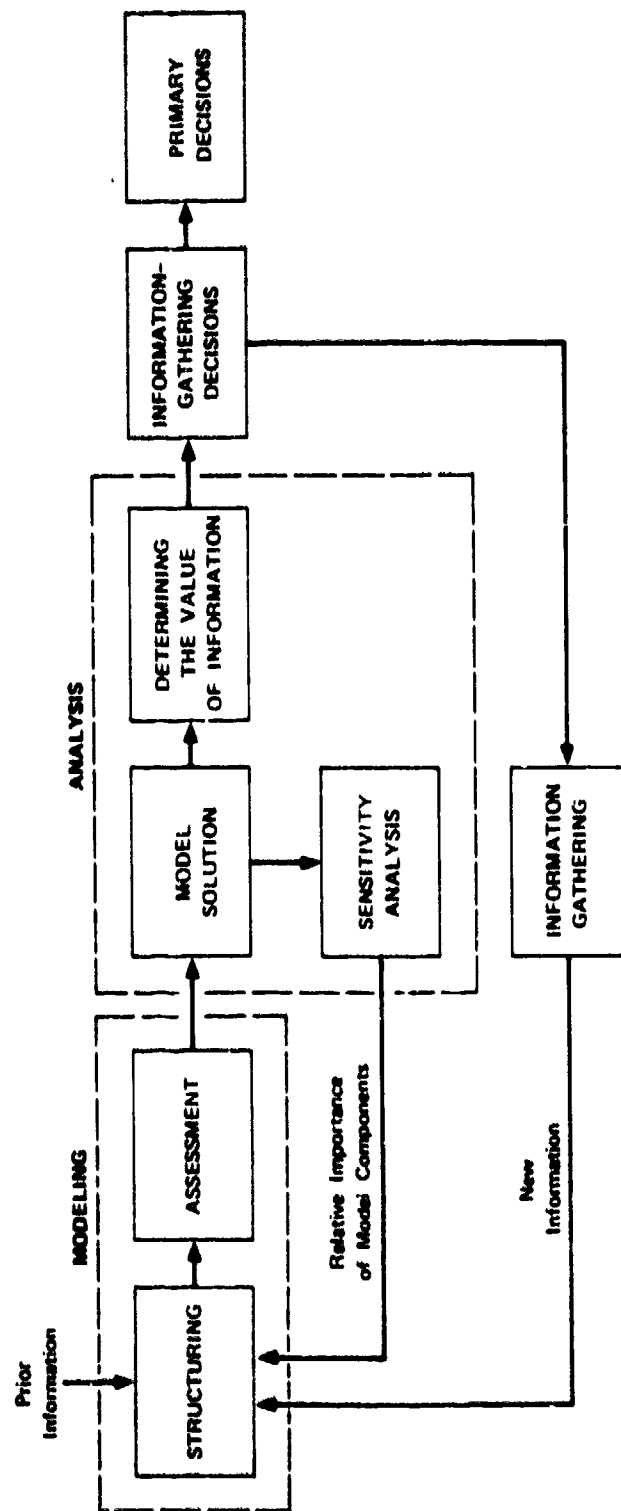


FIGURE 4.1 THE PROCESS OF MODELING AND ANALYZING DECISIONS

## 1. Structuring

During the structuring process, a decision model is defined in skeletal form. The decision maker's alternatives, the possible outcomes that could result from his actions, the external factors affecting the outcome, and the linkages among the model variables are all identified.

Usually the most critical assumptions contained in a model--assumptions that define the scope and nature of the model--are made when its structure is defined. However, there are no rigid rules for structuring decision models. Compared to methods of assessment or analysis, the structuring process has been the object of relatively little research and theoretical development. The discussion that follows outlines typical approaches to structuring decision models, but there are probably as many variations to these approaches as there are analysts and decision makers.

The first step in developing and appropriate model structure is making sure that the right decision is being analyzed. The need for a decision may be recognized, but if the problem is defined so that major alternatives or important elements of the environment are not considered, a comprehensive analysis will never be possible. A decision situation can be put in proper context by identifying the decision makers involved and the resources that each has to allocate. This simple step can quickly eliminate facets of the problem about which a decision maker is concerned but over which he has no control. It is quite possible that the initial perception of a decision situation is unnecessarily narrow or restricted. Probing the constraints placed on a decision maker and exploring his relationship with others who influence or are affected by his decision often changes the perceived nature of the decision. This initial exploration of a decision is carried out qualitatively and intuitively, but it forms a conceptual framework for a quantitative model.

Once a decision has been examined qualitatively and its scope has been properly specified, the system variables that characterize the most important elements of the decision environment are defined. Relatively global system variables that compactly represent the environment are generally defined first, and then are decomposed into more specific variables

as the analysis progresses. Typically the definition of system variables begins with specification of the outcome variables, and then state variables are defined as they are needed to help determine the outcomes of importance to the decision maker. This process leads to questions like, "What would you most like to know in order to specify your state of information about a particular outcome variable?"

The process of defining state variables by working backwards from the outcomes can lead to an unnecessarily large model--one that contains too many system variables--if steps are not taken to focus the analysis on the choices available to the decision maker. Thus in defining the state variables that would be most useful in determining the outcomes, it may be necessary to concentrate on those that are influenced by or define the effects of a decision. In this case the question might be asked, "Which quantities, over which you can exert some influence, would you most like to know in order to specify your state of information about an outcome?"

Obviously the process of working backwards from a variable to others that influence it can be applied to any state variable, not just outcomes. For instance, having determined that the equivalent annual cost of a large piece of equipment is determined by its operating cost, maintenance cost, and depreciated purchase cost, it then may be useful to specify the maintenance cost as a function of the mean time to failure, the repair time, and opportunity cost of maintenance personnel. However, the extent to which the model should be refined and expanded can be guided by a preliminary analysis of the model as it is being developed. (See the discussion of sensitivity analysis later in this chapter.)

Another procedure used to specify system variables is to identify a sequence of events that could lead from a decision to an outcome. The development of a scenario shows only one of the possible consequences of a decision, but a careful definition of a scenario usually exposes the major variables that might lead to different outcomes. The process of using scenarios to reveal the system variables that should be included in a decision model is most useful in situations where it is possible to establish the sequence of events over time.

It is often the case that these two processes--working backwards from the outcome variables to system variables upon which they depend and specifying of sequences of events over time--lead to definitions of the same system variable, but to different sequences of information states about the variables. Specifying system variables that help determine the values of other variables implicitly defines the order in which the state of information about the variables should be assessed. On the other hand, specifying the sequence of events that could follow a decision defines the order in which information about the variables will actually be revealed. As discussed in the preceding section, a decision model may require two sequences of information states: one for assessment and one for model solution. A fully developed decision model must specify both sequences.

The sequence of information states chosen for assessment strongly influences the ease and feasibility of assessing the existing state of information about the decision environment. Fortunately an analyst usually is free to specify this sequence. However, the sequence used to solve a model is determined by the decision environment. This sequence is one of the basic assumptions that determines the realism of a model. This solution sequence may be especially difficult to determine in situations where events can occur in several possible sequences. For example, a model concerned with the U.S. response to intervention by any of several major powers in a regional conflict between two small countries may strongly depend on the order in which it is assumed the major powers will become involved. If either sequence of information states is not obvious from the decision environment, it may be necessary to test the sensitivity of both the model's inputs and outputs to different sequences.

The interdependencies among system variables are often identified at the same time that the variables are defined, but this sort of implicit definition of model linkages may lead to an inaccurate or insoluble model. It is important to include significant interdependencies among system variables in the model, but since the number of possible interdependencies grows very rapidly as the number of system variables is increased, numerous independence assumptions are required in all but the most trivial models. An important part of defining the structure of a decision model is keeping



track of the linkages among variables and the resulting flows of information through the model. Independence assumptions that individually make sense can easily produce unrealistic and unanticipated results if their combined effect is to isolate or overemphasize portions of the model. Several methods for visualizing the linkages among system variables are discussed in the following section.

Most of the structure of a decision model is developed to relate the decision maker's choices and external influences to the outcome variables. However, the designer of a decision model must also specify the structure of the preference model used to evaluate the outcomes. This is a relatively simple process since there are only a few basic types of preference models, as described in the preceding section. To specify the structure of a preference model one must decide whether the preference model will include a separate value model, risk-preference model, and time-preference model; and which procedures will be used to elicit the desired value measures and utilities (such as the multi-attribute approach, definition of tradeoffs, paired comparisons, etc.). The elicitation procedures must be compatible with the type of preference models chosen. For instance, elicitation procedures that produce ordinal quantities should not be used with preference models that require cardinal quantities.

As the structure of a decision model is defined, the model components can be tested and modified. The process of testing model components relies on the judgment of the decision maker and his advisors, and preliminary analyses of an evolving model. Tests of model components can be placed in three categories: meaningfulness, completeness, and appropriate level of detail. These tests are applied repeatedly as a decision model is developed.

Testing a model component for meaningfulness is equivalent to asking, "Does the proposed presentation make sense to the decision maker?" If the answer to this question is negative, then one or more of the model components need to be reformulated to make them correspond to the decision maker's perception of the situation.

Testing a model component for completeness means checking to see whether some important aspect of a decision problem has been overlooked in the current version of the model. A typical question corresponding to this test is, "Does the model include all of the alternatives available to the decision maker?" If the answer to this sort of question is negative, then the type of component under consideration (in this case, the decision variables) must be expanded to include the missing elements of the problem.

Testing a model component for appropriate size means checking the level of detail at which some aspects of a decision problem are represented in a model and the amount of assessment and computational effort implied by such a representation. A typical question corresponding to this test is, "Does the importance of this aspect of the problem justify the detail with which it is modeled?" If the answer is no, the model can be contracted by eliminating or aggregating noncritical components.

The structure of a decision model is limited by the extent to which a decision maker can anticipate changes in his view of the environment resulting from new information. The effect of changes in the decision maker's state of information about the appropriateness of a particular model structure becomes especially important in an attempt to model a decision that must be made some time in the future. While a current analysis of some future decision may clearly indicate a preferred alternative, a new and unanticipated perception of the problem structure may lead to the selection of another alternative when the decision is finally made. This possibility makes it undesirable to analyze a decision completely until it is possible to make an immediate choice. Instead the analysis of a future decision should be based on a structure with enough generality to accommodate possible changes in the decision maker's perception of the problem. However, the fact that every possible future state of information and its effect on the structure of a model cannot be foreseen does not mean that comprehensive decision models are inappropriate for current decisions. Decisions can only be made with whatever information is available at the time; a model simply serves to explore the implications of that state of information.

## 2. Assessment

During the assessment process the existing state of information about components of a decision model is quantified. The values of probability distributions for state variables are specified, and the exact nature of the transformations linking system variables is defined. In addition, the decision maker's preferences among possible outcomes, the time at which they might occur, and the risks associated with each alternative are specified.

Fortunately the process of assessing states of information and preferences has been the object of considerable research. While there is not universal agreement on how these assessments should be made, the various procedures that have been developed are characterized more by their similarities than their differences. Typical assessment procedures are discussed below.

### a. Assessing State Variables

The manner in which information about a state variable is assessed depends on how important the variable is to the output of decision model. Information about critical state variables, those to which the outcomes are highly sensitive, is usually assessed in the form of probability distributions, while less important state variables are usually specified deterministically. Development of procedures for assessing state variables has concentrated on various types of probability encoding.

Most probability encoding procedures have dealt with the question of eliciting information from an individual subject. This form of encoding must overcome the cognitive and motivational biases that distort the subjective information supplied by an individual. Cognitive biases are a systematic distortion of subjective estimates caused by the way one thinks about uncertainty. For example, a response may be biased toward the most recent piece of information simply because that information is easiest to recall. Motivational biases are distortions in an individual's subjective judgment caused by his perceived system of personal reward for various responses. For example, an individual may want to bias his response

because he perceives that his subsequent performance will be evaluated by a comparison of his response to the actual outcome.

Techniques for eliciting information may be classified according to the type of questions asked. Probability methods require the subject to assess the probability associated with a particular value of a state variable. Value methods require the subject to assess the value that corresponds to a particular probability. Mixed probability-value methods require the subject to assess values on both the outcome and probability scales; the subject essentially describes points on the cumulative distribution of an uncertain quantity.

These encoding methods can be used in either a direct or an indirect response mode. In the direct mode, the subject is asked questions that require numbers as answers. The answers can be given in the form of probabilities (or equivalently in the form of odds) or values. In the indirect response mode, the subject is asked to choose between two or more uncertain situations. The probabilities characterizing the situations are adjusted until he is indifferent, and the point at which he is indifferent is translated into a probability or value assignment. The indirect response mode is typically used with a reference process, where the subject is asked to compare some aspect of an uncertain quantity to a reference process such as the toss of a fair coin or the spin of a wheel of fortune.

Which encoding method is most appropriate depends on the type of uncertainty being assessed. For instance, if there are only a few possible outcomes for an uncertain quantity, a method that requires the subject to divide the range of possible values for a state variable into a number of intervals may not be appropriate. The choice of encoding method should be based on the characteristics of the uncertain quantity, its importance to the modeling effort, and the personal preferences of the person supplying the information.

The encoding process consists of five distinct phases: the motivational phase, structuring phase, conditioning phase, encoding phase and the verification phase. In the motivational phase, the analyst must explore the subject's motivational biases and attempt to eliminate or

compensate for them. In the structuring or definition phase, the subject and analyst must reach an agreement on the exact definition of the uncertain quantity being considered. The conditioning phase is directed toward finding out how the subject goes about making his probability assignments and heading off any biases that might surface during the encoding process. Once the uncertainty has been well defined and the subject's cognitive and motivational biases have been explored, the process enters the encoding phase and utilizes one of the encoding methods discussed above. In the verification phase the encoded information is subjected to a number of consistency checks to see if it truly represents the subject's beliefs.

Encoding the probability of rare events--events that have a very small probability of occurrence--is difficult for two reasons. First the subject is asked to assess the probability of an event with which he has, by definition, little experience. Secondly, he may have to distinguish among very small probabilities (for example, a probability of one in one thousand as compared to a probability of one in ten thousand).

Procedures have been proposed for overcoming these difficulties. One such procedure is to relate the rare event to other uncertain events with which the subject is more familiar; another is to describe the event in terms of its component parts or as the result of a sequence of other events. The purpose of these procedures is threefold: to gather all of the information relevant to the event that is within the subject's command, to see if the event can be redefined in terms that no longer involve small probabilities, and to produce lists of related events that can be used as reference points in getting the subject to make relative statements about their uncertainty.

The subject is asked to make relative judgments about the occurrence of a number of different events before dealing with very small probability numbers. Subjects can often say whether one rare event is more or less likely than another rare event, even though they would find it difficult to assign a numerical probability to either event. By doing so the subject can bound the probability associated with a rare event, and then narrow the bounds as required for the modeling effort.

The assessment of state variables becomes more complicated when the judgment of a group of people is quantified. There are many decision problems for which more than one expert is available to supply subjective estimates of uncertain quantities, and there can be considerable differences of opinion among the various experts. Group encoding procedures are used to aggregate the estimates of several experts into one probability distribution that can be used in a decision analysis.

Group encoding is currently carried out with several subjects, whose judgments may have been encoded individually prior to the group encoding session, and an analyst who monitors the flow of information among the subjects to see that all the relevant opinions are discussed and incorporated into group consensus. Another individual may have the responsibility of reviewing the group consensus and dissenting opinions, and deciding on the final probability distribution that should be used in the analysis. This individual is typically the head of the organization that has employed the services of the group of experts.

Much of the research associated with group probability encoding has been concerned with the way in which information should be exchanged among the individuals supplying subjective estimates. At one extreme, it has been suggested that participants in the group encoding process remain anonymous, and that they send each other their subjective estimates but not the logic behind the estimates. This procedure has been called the Delphi technique; it is designed to reduce the influence of group members with dominating personalities and to encourage each member of the group to reach his or her own opinion. Alternative procedures allow the participants to meet face-to-face and exchange information freely. These procedures are often more successful in producing a group consensus than is the Delphi technique, in part because each member of the group can learn from the additional information supplied by others.

b. Assessing Dependencies

After the structuring process has identified the linkages among system variables, the manner in which the variables depend on each other is usually assessed in two steps. First the dependencies are assumed

to have particular forms, and then they are assessed quantitatively. For instance, it might be assumed that the cost of maintaining a piece of equipment grows exponentially, so the cost in any year can be determined by multiplying the previous year's cost by a certain growth factor. With this assumption, the quantitative assessment reduces to one of specifying the growth factor.

In determining the nature of a decision model, the assumptions made about the form of interdependencies among variables are at least as important as a quantitative assessment of the dependencies themselves. Yet most of the existing procedures for assessing dependencies focus on the quantitative assessment of relationships whose form has already been determined. For example, regression analysis (linear least-squares curve fitting) provides a means for deriving the parameters of a transformation from existing data if one is willing to assume that the transformation has a particular linear form. While assumptions about the form of variable dependencies may be necessary to simplify the task of assessing them, caution must be exercised to avoid assumptions that do not approximate reality. No amount of quantitative assessment will produce an acceptable approximation of a relationship among elements of a decision problem if the relationship on which the assessment is based has an inappropriate form.

One of the basic assumptions that must be made about the form of relationships among system variables is whether they should be represented deterministically or probabilistically. By assuming that a dependency is probabilistic, one can include the effects of several possible deterministic transformations in the same mode. This makes it somewhat easier to accept the assumptions inherent in the form of the transformation. However, the use of probabilistic relationships can greatly increase the difficulty associated with quantitatively specifying the dependency. To overcome this difficulty it may be necessary to couple the use of probabilistic relationships with assumptions of independence, especially in models where state variables are aleatory and the transformations among them are described by conditional probability distributions.

For example, consider a model in which there are four aleatory state variables, each of which has three possible values. If the variables are all independent, the decision maker's state of information about them can be specified with 12 probabilities. On the other hand, if they are totally dependent, it will be necessary to assess at least 81 probabilities, most of them defining probabilistic dependencies.

Sometimes it is easier to assess the uncertainties associated with dependent aleatory state variables as joint probabilities, while in other situations it is easier to assess the same information in terms of conditional probabilities. A joint probability distribution specifies the likelihood of every possible combination of values for the state variables. Conditional probability distributions specify the probability that some subset of state variables have certain values given the values of the remaining variables. Either set of probability distributions can be derived from the other; the decision maker can supply whichever is easier to assess.

In the example of four aleatory state variables, each of which has three possible values, there are 81 joint probabilities, one for every possible combination of values. However, 117 conditional probabilities, plus three additional probabilities specifying the state of information about one state variable, are required to specify the same information if the variables are completely dependent. In this case the number of conditional probabilities needed to specify the dependencies exceeds the number of joint probabilities, which means the joint probabilities are probably easier to assess. However, if it is assumed that some of the state variables are independent, the number of probability assessments may be minimized by eliciting conditional, rather than joint, probabilities. For instance, if each state variable in the previous example depends on only one other variable, only 27 conditional probabilities are required to specify the dependencies. In this case it is probably preferable to assess conditional probabilities.

However, the choice of assessing either conditional or joint probability distributions to represent probabilistic dependencies can depend on factors other than the required numbers of probability



assessments. Conditional probabilities are often difficult to assess when the state variables represent events that can occur in any order. When a unique sequence of events does not exist, it may not be obvious how the dependencies in a model should be defined or assessed. Assessing conditional probabilities in different orders can lead to conflicting sets of probabilities. Another problem with assessing dependencies in terms of conditional probabilities is that the subject may anchor on a probability given one state of information, and then shift his probability estimates by an insufficient amount when the state of information is changed. To overcome these difficulties it may be necessary to assess the dependencies among aleatory state variables in several different orders to expose and correct inconsistencies, or assess the dependencies in terms of joint probability distributions.

c. Assessing Preferences

The process of assessing preferences follows from the structure that has been defined for the preference model. If distinct value, risk, and time preference models have been specified in the model structure, then each type of preference is assessed separately. Otherwise, all types of preferences are assessed together. Several assessment techniques have been developed for use with various sorts of preference models. Each technique requires the decision maker to make representative choices, and then this information is used to infer his preferences over a certain range of outcomes. The techniques differ in the way that the representative choices are presented to the decision maker, and the manner in which his responses are processed to specify the preference model.

One way to encode preferences is to ask a decision maker to choose between two possible combinations or outcomes, or, if his risk attitude is being encoded at the same time, between probability distributions over the outcomes. This is called the "paired comparison" method. If the decision maker indicates a preference, the outcomes (or probability distributions over outcomes) are modified to find a set of outcomes (or distributions) between which he is indifferent. If the outcomes are

deterministic, the decision maker's preferences can be represented by indifference curves in the space spanned by the outcome variables. If the decision maker is offered choices among probability distributions over the outcomes, it is generally not possible to draw indifference curves, and a large number of inquiries may be necessary to establish his preferences. Thus paired comparisons are usually made between combinations of outcomes that are known with certainty, which is equivalent to splitting the preference model into a separate value model and risk preference model.

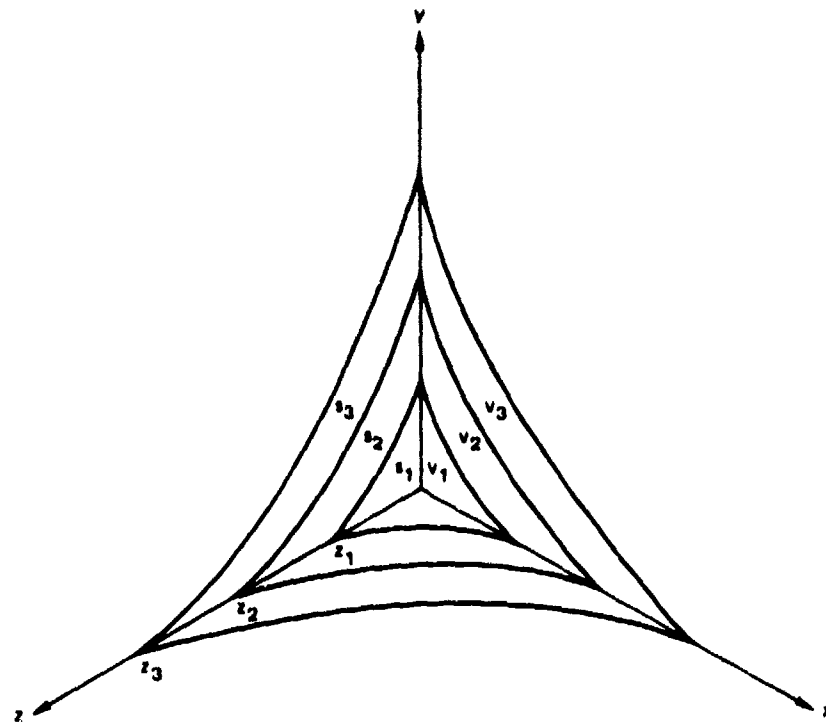
The paired comparison method can be generalized slightly by asking the decision maker to rank several possible combinations of outcomes in terms of his preferences. This amounts to making a number of paired comparisons at the same time. In any case, when this method is applied to deterministic outcomes, it produces an ordinal value measure since the decision maker is never asked to say how much he likes one combination of outcomes more than another. If the paired comparison method is applied to probabilistic outcomes, it produces a cardinal utility measure with the properties discussed in Chapter III.

Another way to encode preferences is to ask a decision maker to specify tradeoffs among various outcome variables. This means that he must specify how much he would be willing to change one outcome variable in an undesirable direction in order to make a desirable change in another outcome variable. By specifying such a tradeoff, a decision maker defines two sets of outcomes between which he is indifferent: the original set of outcomes and a new set in which two of the outcomes have been changed. This method of assessing preferences is almost always used with deterministic combinations of outcomes, which means that a separate risk preference model is required. It also means that the tradeoff method can usually be used to develop indifference curves. Since the decision maker is asked to specify only deterministic combinations of outcomes among which he is indifferent and not how much more one combination is worth than another, the tradeoff method produces an ordinal value measure.


The value measures produced when the paired comparison and tradeoff methods are applied to deterministic combinations of outcomes are often interpreted as cardinal quantities, even though both methods produce ordinal quantities. As explained below, this interpretation can be justified if the value measure is defined in terms of an appropriate outcome.

At the completion of the assessment process the decision maker's preferences are defined as a ranking of all possible combinations of outcomes. This ranking is usually extrapolated from his choices among representative outcome combinations. The preference ranking can be represented graphically as indifference surfaces or curves, as shown in Figure 4.2(a). This figure shows some of the indifference surfaces that might be assessed for three outcome variables  $x$ ,  $y$ , and  $z$ . The assessed value model maps the combinations of outcomes represented by the three indifference surfaces--  $S_1$ ,  $S_2$ , and  $S_3$  --into three levels of the value measure--  $v_1$ ,  $v_2$ , and  $v_3$ . If the combinations of outcomes corresponding to the  $S_3$  indifference surface are preferred to those corresponding to the  $S_2$  indifference surface then  $v_3$  exceeds  $v_2$ . The amount by which  $v_3$  exceeds  $v_2$  is irrelevant for defining preferences. In fact, any positive, monotonic transformation could be applied to the value measure and it would still describe the same preferences and indifference surfaces. The transformation would change the values of  $v_1$ ,  $v_2$ , and  $v_3$  but not their ranking; hence, a preferred combination of outcomes would still have a higher value measure.

There are many possible ways to define the value measure, since it is arbitrary within a positive, monotonic transformation. One way the value measure could be defined is in terms of the outcomes. For instance, if for any combination of outcomes, there exists another combination whose outcomes are all zero except for one particular outcome such that the decision maker is indifferent between the two combinations, then the value measure can be defined as that non-zero outcome. This is shown in Figure 4.2(a). The three indifference surfaces--  $S_1$ ,  $S_2$ , and  $S_3$  --each contain one combination of outcomes where  $x$  and  $y$  are zero. These outcome sets correspond to the points where the indifference surfaces



(a) INDIFFERENCE SURFACES FOR THREE OUTCOME VARIABLES



(b) INDIFFERENCE CURVES WHEN  $z$  EQUALS ZERO

FIGURE 4.2 INDIFFERENCE CURVES ON SURFACES

intersect the  $z$  axis, and the values of  $z$  for these outcome sets are  $z_1$ ,  $z_2$ , and  $z_3$ . These values of  $z$  can be used as the value measure since preferred combinations of outcomes have larger values of  $z$ .

If the value measure is defined by a cardinal outcome variable, then the value measure can be treated as a cardinal quantity. For instance, if  $z$  in Figure 4.2(a) is incremental wealth (measured in dollars, for instance) and  $z_3$  is three times  $z_1$ , then it is possible to interpret a set of outcomes represented by a point in the indifference surface  $s_3$  as being worth three times as much as a set of outcomes represented by a point in  $s_1$ . In particular, the value measure for sets of outcomes involving only  $x$  and  $y$  can be placed on a cardinal scale by associating the appropriate value of  $z$  (wealth in this case) with each  $x$ - $y$  pair. Preferences among outcomes involving only  $x$  and  $y$  can be represented as indifference curves, as shown in Figure 4.2(b). These curves are equivalent to those shown in Figure 4.1(a) where the indifference surfaces-- $S_1$ ,  $S_2$ , and  $S_3$ , and the value measures for these indifference curves are  $z_1$ ,  $z_2$ , and  $z_3$ .

Often indifference curves such as those shown in Figure 4.2(b) are assessed together with an equivalent level of some other outcome, but without an assessment of the complete indifference surfaces shown in Figure 4.2(a). Typically the decision maker is asked, "How much of one particular outcome would you accept in place of a given set of outcomes (or in place of any other outcome set that you find equally desirable)?" In this way, a situation with non-monetary outcomes might be evaluated in monetary terms. However it is not necessary for the value measure to be defined in terms of money in order for it to be considered a cardinal quantity. Any well-defined, cardinal, quantitative outcome--such as damage to the enemy or lives lost--will suffice.

If the value measure is not defined in terms of a cardinal outcome variable, it cannot be treated as a cardinal quantity. For example it is not appropriate to ask a decision maker to state how much more he likes one combination of outcomes than another, or to specify the relative "value" or "utility" of two sets of outcomes. In this case the decision maker is asked to specify a cardinal quantity on an ordinal scale. The response to such a request is not well defined since

it cannot be known which scale the decision maker used to calibrate his preferences.

A third way of assessing preferences is the multi-attribute approach discussed in Chapter III. This method produces an individual value measure for each of the outcomes (attributes). The individual value measures must be assessed as cardinal quantities so that they can be combined to form a single, composite value measure. This means that each of the value measures must be defined in terms of a cardinal outcome variable. It is not necessary that all of the individual value measures be defined in terms of the same outcome variable, but the process of combining the individual value measures requires fewer assumptions if this is the case. In any case, asking a decision maker to define his preferences with respect to each of the outcomes (attributes) in terms of "value" or "utility" produces a set of ordinal quantities that cannot be algebraically combined to produce a composite value measure.

It is often the case that several individuals with different preferences are involved in a decision. A complete analysis of such a situation requires consideration of the alternatives open to each individual to influence the primary decision. Game theory and social choice theory are appropriate for such an analysis. However, when there is an identifiable decision maker, it is clear that his preferences will ultimately be used to make the decision. The desires and possible actions of others may figure into his preferences, but a decision maker will eventually have to choose an alternative that he thinks is best. When it is not clear who is the decision maker, it may still be possible to analyze a decision situation and show that the range of preferences relevant to the decision all lead to the same preferred alternative. At a minimum, an analysis of a situation in which there is no single decision maker can focus debate on those preferences issues that lead to the selection of different alternatives.

If it is necessary to assess a separate risk preference model, the decision maker is asked to make several representative choices among probability distributions over the value measure. Typically the

distributions are simple, and the possible levels of the value measure are chosen so that the utility associated with all but one of them is known. The distributions are altered until the decision maker is indifferent between them, and the unknown utility is determined by equating the expected utilities associated with the two distributions. Often it is assumed that the decision maker's risk aversion is constant over a given range of value measures. In this case the decision maker's choices among reference lotteries are used to determine and verify the constant level of risk aversion. The most common assumption made with respect to risk preference is that the decision maker is risk neutral. This means that he is assumed to choose the value distributions with the highest expected value.

Methods currently used for assessing time preference are clumsy at best. Typically it is assumed that the decision maker's time preference can be represented by a constant discount factor, which is then treated as a system variable. More complicated and realistic time preferences are usually assessed by considering the possible occurrences of an event in different time periods to be different outcomes, and then using the methods discussed above to assess preferences among outcomes. This amounts to including time preference in the value model and not constructing a separate time preference model. However, recent results [1] indicate that practical methods for assessing a separate time preference model are possible, if certain assumptions are made about the form of a decision maker's aversion to variations in the value measure over time.

#### B. Analysis

As shown in Figure 4.1, there are three steps in analyzing a decision model: the solution of the model, sensitivity analysis, and determination of the value of information. The solution of the model to find the recommended alternative and the corresponding distribution of possible outcomes is primarily analytical; it forms the basis for the other two. Sensitivity analysis is used to determine the relative importance of model

components and thus guide efforts to revise and improve the model. Determining the value of information provides the basis for decisions to gather new information, which can be used to update or restructure the model.

#### 1. Model Solution and Decision Optimization

There are several ways to solve any decision model. Which method is most efficient depends on the nature of the model and the accuracy required in the solution. Even if a model has been constructed so that it can be solved using a particular technique, there may be other solution techniques that are equally if not more efficient. For example, a decision tree can be solved using direct processing of the tree structure, proximal analysis, Monte Carlo techniques, dynamic programming, and linear programming. Some of these solution techniques are relatively inefficient, but they can provide insights not produced by other techniques. For instance, determining the second-best decision strategy in a complicated decision tree is difficult using procedures that process decision trees directly. However, a linear programming solution of the same problem easily produces the desired result. Switching from one solution technique to another may require a reformulation of the decision model, but the effort involved is often justified by the savings in computational effort required to solve the model.

The direct solution of decision trees, one of the basic types of models used in decision analysis, is discussed in any elementary text on decision analysis [13]. This solution technique, sometimes called "rolling back" a decision tree, is usually adequate for simple decision problems. However, a characteristic of decision trees is that they tend to grow very rapidly as the size of the model is expanded. At present, it is not uncommon for analysts to utilize decision trees containing thousands, or even millions, of nodes. Even with the use of automated decision aids, straightforward procedures for analyzing decision trees can become time-consuming and expensive. A method of overcoming this problem, by eliminating redundant portions of a decision tree, is described in Chapter V.



Another way to analyze large probabilistic models is to use approximate methods for solving them. Perhaps the most common of these is the Monte Carlo method, which involves repeatedly sampling the probability distributions for the aleatory state variables and then solving the remainder of the decision model deterministically. If this process is carried out many times for each of the alternatives available, the deterministic model solutions will collectively approximate the distribution of outcomes associated with each alternative. The Monte Carlo method has the advantage that it does not require probabilistic processing of the model, and the disadvantage that it may require a very large number of deterministic model solutions before an adequate approximation of the distribution of outcomes can be obtained. Furthermore, if there are a large number of alternatives (which is often the case when there are several decision variables), a prohibitively large number of deterministic model solutions may be required to find the distribution of outcomes associated with each alternative.

Approximate methods have also been developed for analytically processing probabilistic models [6]. These methods, called proximal analysis, make it possible to determine approximate values for the mean and higher moments of an uncertain outcome in terms of the moments of the distributions for the aleatory state variables, without completely solving the model for the probability distribution over the outcomes. The accuracy of the approximations depends on the nature of the transformations from the aleatory state variables to the outcome variables. The advantage of proximal analysis is that it approximates moments of the distribution of outcomes without repeated deterministic solutions or a complete probabilistic solution of the decision model.

A complete description of all of the methods available for analyzing decision models, including probabilistic models, is beyond the scope of this report. However the various analytical methods are well documented, and are probably the most clearly defined elements of the decision analysis process [4,12,13]. The importance of these methods for the overall process of modeling and analyzing decisions is that they determine the cost, and even the feasibility, of using the model developed

by the other steps to gain insight into a decision situation. The availability of an efficient solution technique can justify the development of a complex decision model, as long as the model is an adequate representation of the decision environment and not just one that is compatible with the solution technique. Thus solution techniques are a guide to the modeling process since they determine whether the benefits to be derived from the model will exceed the cost of analyzing it.

## 2. Sensitivity Analysis

One of the basic tools for analyzing decision models is sensitivity analysis. By determining the sensitivity of one model component to changes in another, one can gain insight into the relative importance of various components and the effect that modeling assumptions have on the output of the model.

There are a number of ways to measure the sensitivity of one parameter to another. If both parameters are known with certainty, deterministic sensitivity analysis can be used. The most common way to carry out this form of analysis is to vary one of the parameters over a specified range, and plot the corresponding values of the second parameter as a function of the first parameter. The resulting graph, showing the pairs of parameter values that fit the model, describes the global sensitivity of the parameters to each other. The analysis is global in the sense that the relationship between the two parameters is described over a given range. Alternatively, the local sensitivity of one parameter to another at a given operating point can be computed by determining the rate at which small changes in one parameter cause the other parameter to change. Local sensitivities are measured in terms of partial derivatives. All sensitivity analyses can be classified as either global or local.

Another way to classify sensitivity analyses is by whether they produce open-loop or closed-loop sensitivities. A closed-loop sensitivity is one in which the decisions embedded in the model are reoptimized whenever one of the parameters is changed. An open-loop sensitivity, which is usually easier to compute than a closed-loop sensitivity, is one in which the decisions in the model are fixed throughout the analysis even

though changing one of the parameter values over its range might change the optimum decision.

For uncertain model parameters, a probabilistic sensitivity analysis is possible. Typically this analysis is carried out by holding an aleatory state variable constant at each of its possible levels and observing the effect on the output distribution. However, there are several other ways to conduct a probabilistic sensitivity analysis. In general, a probabilistic sensitivity analysis involves modifying the probability distribution for one state variable and observing of the effect on the distribution for other state variables. The procedures for probabilistic sensitivity analysis are not as well defined as those for deterministic sensitivity analysis because there is a variety of ways in which one can modify a probability distribution. For example, in one analysis it might be desirable to change the variance of one distribution while holding its mean constant; in another analysis it might be preferable to vary algebraic coefficients in the expression for the probability distribution. As with deterministic sensitivity analyses, probabilistic sensitivity analyses can be either global or local, and either open-loop or closed-loop.

Joint sensitivity, the effect of varying several model components simultaneously, is more difficult to calculate and visualize. The problem with calculating joint sensitivities is that there can be very many ways to vary several variables simultaneously. Furthermore, visualizing global joint sensitivities requires plotting the dependent variable in several dimensions. One way to avoid these problems is to determine local joint sensitivities in terms of gradients. Gradients can be represented as  $n$ -dimensional vectors and are much easier to comprehend than  $n$ -dimensional functions. Joint sensitivities to probabilistic quantities are even harder to visualize than joint deterministic sensitivities, and are therefore rarely used.

One of the primary reasons for carrying out a sensitivity analysis is to manage the growth of a decision model. Decision models are commonly designed from the "top down." A few aggregate system variables are defined first, and then they are defined in terms of increasingly

detailed and specific variables. For example the outcome of a public sector decision might be a measure of social profit. However, directly estimating this measure for each alternative is equivalent to selecting an alternative by intuition. To model the situation, the measure of social profit can be divided into benefits and costs. Both of these variables can be further expanded by defining them in terms of more specific quantities until a level is reached where experts are available to estimate the inputs to the model. The process of expanding the model is continued until it is no longer economic to do so.

Figure 4.3 shows how top-down modeling and sensitivity analysis affect the size of a model. At the start of the modeling process, intuition




FIGURE 4.3 INITIAL MODEL PROGRESSION IN DECISION ANALYSIS

and direct assessment are used to predict the outcomes and some inferior alternatives can be discarded without formal analysis. As the modeling process progresses, the size and complexity of the model increase. Then sensitivity analysis is used to eliminate relatively unimportant variables, and the size of the model contracts.

The process of modeling and analyzing a decision rarely ends after the first application of sensitivity analysis. Instead, the information provided by a sensitivity analysis is used to identify areas where additional modeling is required. This leads to another expansion of the model as shown in Figure 4.4. The result is that the size of the model alternately expands and contracts, with the average model size increasing gradually as important variables are defined, tested with sensitivity analysis and retained.

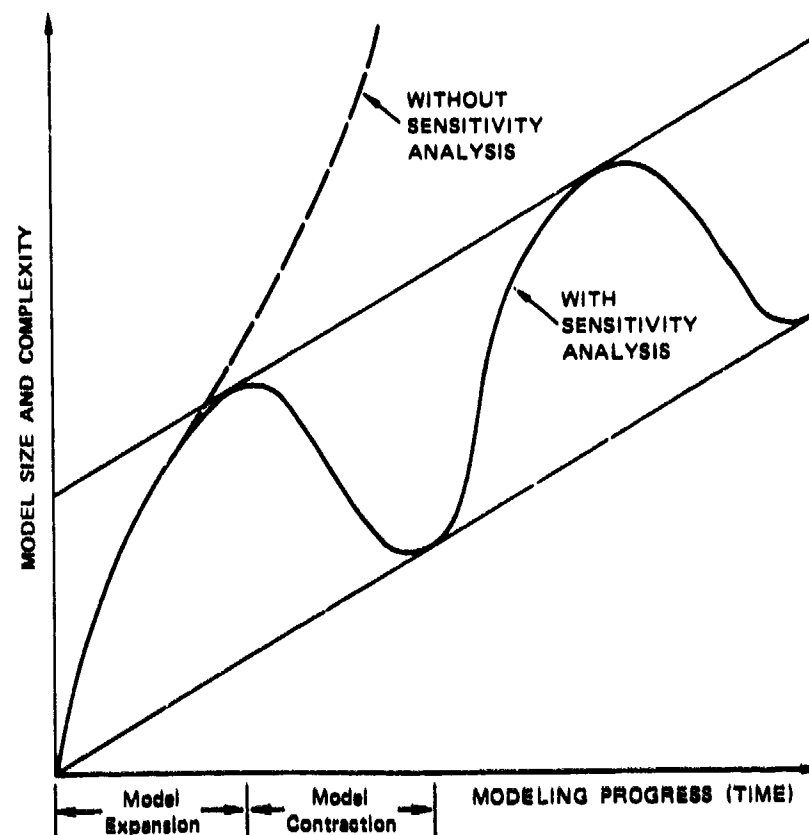


FIGURE 4.4 MODEL PROGRESSION IN DECISION ANALYSIS

Each time the model expands and contracts the decision analysis process has proceeded through the steps of structuring, assessment, model solution, and sensitivity analysis, as shown in Figure 4.1. Eliminating unnecessary model components is one of the important characteristics of decision analysis. If a decision model were developed without discarding relatively unimportant variables, the size of the model would grow rapidly as shown by the dashed line in Figure 4.4.

### 3. Determining the Value of Information

Sensitivity analysis is not the only tool for evaluating the relative worth of various model components. Determining the value of information can also supply a measure of the importance of state variables, in addition to guiding efforts to gather new information.

One of the advantages of a probabilistic model is that it can be used to calculate the value of information. Information changes the uncertainty inherent in a decision, and this change can be incorporated in a probabilistic model since uncertainty is represented explicitly. In situations where the decision maker's attitude toward a risky situation is independent of his wealth, the value of information can be determined by solving a decision model twice--once with the original state of information and once with new information--measuring the difference in the "certain equivalent". The certain equivalent is the deterministic level of the value measure that the decision maker views as equivalent to a probability distribution over the value measure. However, in situations where a decision maker's attitude toward a risky situation depends on his assets, the cost of obtaining information must be included in the decision model. The value of information is equal to the cost such that the decision maker is indifferent between purchasing the information or acting without it.

As shown in Figure 4.1, the possibility of purchasing information presents the decision maker with some secondary, information-gathering decisions in addition to the primary decision. The two types of decisions are often confused. For example, decisions to allocate resources to intelligence activities are information-gathering decisions; they

should be analyzed in the context of other decisions that can be influenced by the information resulting from the intelligence activities. If the primary decisions for which the information is required are not included in the model, it is necessary to assess the value of information directly. Unfortunately, humans are notoriously poor at assessing the value of information. Determining the value of information often produces the most counter-intuitive results of a decision analysis.

The value of information is usually determined for information about aleatory state variables. However, when the dependencies among state variables are uncertain, it is still possible to calculate the value of information about them. In practice the same procedures are used to determine the value of information that would alter the decision maker's perception of a problem and thus require significant structural changes in the corresponding model. If the manner in which the information changes the model can be anticipated, an expanded version of the model that includes this possibility can be used to calculate the value of the information. However, information that changes an individual's perception of a decision is often unanticipated. Unfortunately no amount of modeling will allow one to evaluate unanticipated events or information. Decision models, like decisions, can only be based on what the decision maker knows, or can learn about the situation.

The value of information can be used as a guide to further modeling. If there is reason to believe that further modeling of a state variable will produce a better estimate of the variable, then the modeling may be justified if the value of information about that variable is high relative to the cost of the additional modeling. In any case, new information always leads to an updating of the state of information represented in the model.

Calculating the value of information can be difficult in large models. One way to overcome this problem is to use approximate methods to determine the value of information; these methods are discussed in Appendix B. A more common way to simplify the calculations is to determine the value of perfect information about some of the state variables, and then use the result as an upper bound for the value of realistic,

imperfect information about those variables. The value of perfect information is relatively easy to calculate because it usually does not require updating any of the probabilities in the model.

Even the value of perfect information becomes difficult to calculate if the probability distribution associated with state variables depends on decisions, or if information about a state variable depends on decisions. The problem is that the probability that the information will be of a particular form depends on a decision which, in turn, may depend on the information. This problem can be overcome by breaking up a state variable with decision-dependent information into several separate state variables, one for every alternative upon which the information depends. However this drastically increases the size of a decision model. Furthermore, the new decision-dependent state variables are not necessarily independent of each other, which necessitates assessing the dependencies among them. The problems associated with decision-dependent information are discussed in Appendix C.

Determining the value of information is also more difficult when it is possible to buy several pieces of information--called "observables"--sequentially. Learning one observable can affect not only the primary decision, but also information-gathering decisions with respect to the other observables. As shown in Figure 4.1, the new information contained in the first observable is used to update the model, and then additional information-gathering decisions must be made before the primary decision is made. Since knowing the first observable can help with both information-gathering decisions and primary decisions, the value of the first observable is greater than it should be if it could affect only the primary decision. The prices of the observables affect the decision maker's willingness to buy additional information, and therefore the amount that the value of information is increased by the possibility of buying additional information depends on the prices of all the observables. The manner in which decision problems with sequential information can be solved using computers is discussed in Appendix D.



If a decision problem contains several different decisions, it is necessary to know which decisions will have been made at the time information is received in order to calculate the value of the information. The question of the relative timing of information and decisions is described in terms of the flexibility of the decisions. A flexible decision is one that can be delayed until information can be gathered to guide the decision. The value of a piece of information depends on which decisions are sufficiently flexible to allow their postponement until the information arrives. Thus, it may be necessary to refer to the value of information about a state variable in relation to the flexibility of a given decision variable. The subject of flexibility is discussed in Appendix E.

## V DECISION MODELING CONCEPTS

This section deals with some of the conceptual tools that are used to construct and analyze decision models. These concepts are the practical application of the modeling and analysis process described in Chapter IV and their use defines the characteristics of a decision model as described in Chapter III. It is anticipated that decision modeling concepts such as those described in this section, will form the basis for a system of automated decision aids.

The list of modeling concepts described in this section is not exhaustive. However the discussion starts with basic ideas that are common to almost any decision model, and proceeds to those applied modeling tools that have been useful for a broad range of decision problems. The more applied concepts described here can form the basis for automated decision aids, and as new approaches to modeling decisions are developed, a suitably designed system of decision aids can be expanded to incorporate them.

Decision models are often conceptualized and constructed in both algebraic and graphical terms. Sometimes graphical representations of a model are more useful than algebraic forms, and vice versa. However, there is no single conceptual approach that is "best" for modeling a particular decision. One individual might prefer to visualize a complex decision situation in terms of computational graphs, tree structures, block diagrams, or multi-dimensional graphs, while another person might deal with the same problem using systems of equations, alpha-numeric lists of data, or algebraic computer programs (FORTRAN, ALGOL, LISP, APL, etc.) In fact there is value in using several different conceptual approaches--both algebraic and graphical--for the same problem. Different approaches tend to illuminate different aspects of a decision problem. It is often helpful to switch back and forth between algebraic and graphical representations of the same problem, using whichever concepts facilitate each step in the process of modeling and analyzing decisions.

For example, decision trees are often used to represent complex decision situations. However, decision trees easily can become so large that they cannot be visualized in their entirety. As a result, the structure of a tree may be summarized in a more compact graphical notation, or described by a series of algebraic statements. Since these compact representations are often incomplete, it may be necessary to use several of them to completely specify a large decision tree. Each representation supplies a different perspective on the nature of the model.

#### A. Basic Concepts: Entities and Operators

All models have in common certain fundamental building blocks, called entities and operators. Entities describe the state of the environment or one's perception of the environment, and operators describe the way in which entities are modified and related to each other. Entities can be numbers, arrays, functions of time, strings of alpha-numeric data, algebraic variables, logical variables, complex variables, etc. Operators transform one set of entities into another, and describe dependencies or relationships. The definition of entities and operators can be based on algebra, set theory, probability theory, Boolean algebra, linear algebra, calculus, or any other branch of mathematics or logic.

For example, entity  $E_1$  could be the demand for a particular commodity at some point in time. Entity  $E_2$  could be the equilibrium price of the commodity in a competitive market. Operator  $O$  could transform the demand into an equilibrium market price. An economist might view this transformation in terms of a demand curve like the one shown in Figure 5.1(a). However, the same process can be viewed as an operator (the demand curve) transforming entity  $E_1$  (demand) into entity  $E_2$  (price) as shown in Figure 5.1(b).

##### 1. The Properties of Operators

Operators are defined as functions that produce a value for each output for every possible combination of inputs. Furthermore, operators have a direction; they transform inputs into outputs, but not vice versa. However, it may be possible to reverse the transformation

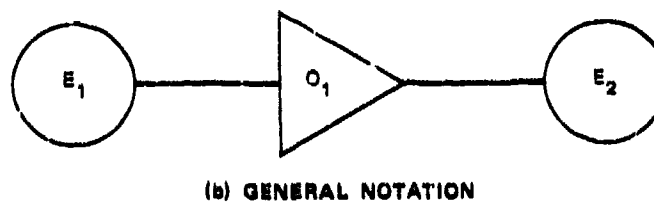
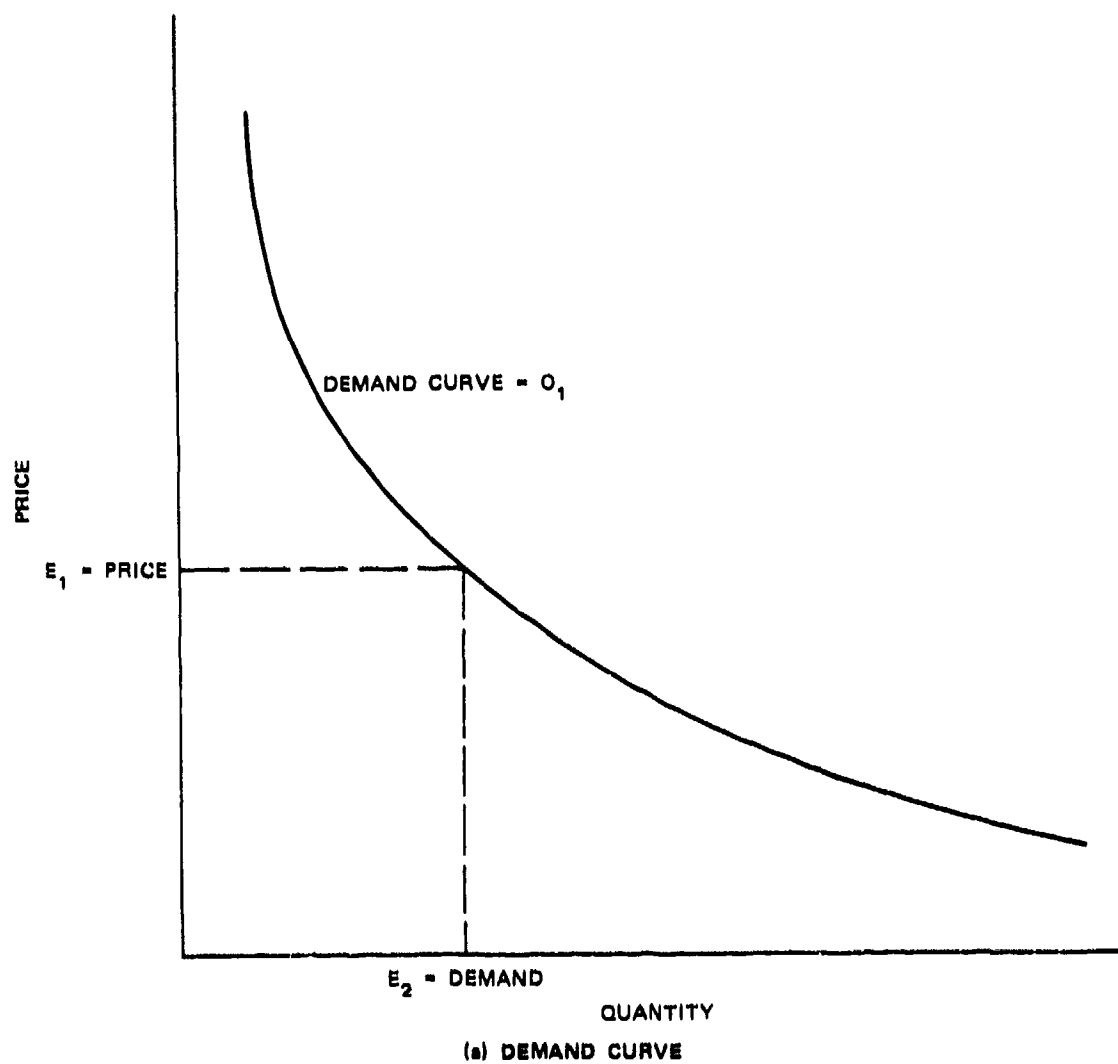


FIGURE 5.1 ILLUSTRATION OF OPERATORS AND ENTITIES USING A DEMAND CURVE

defined by an operator, and thus form a new operator with different inputs and outputs. One implication of the directional nature of operators is that an equation defines several operators. To define a unique

operator, an equation must be accompanied by a specification of its inputs and outputs. For example, the equation  $X = Y + Z$  describes three operators, each of which maps two of the entities--  $X$ ,  $Y$ , and  $Z$  --into a third.

The directional nature of an operator and the requirement that it produce an output make it possible to avoid a number of computational difficulties, such as nonexistent or multiple outputs. However, these requirements may make it difficult to reverse an operator's inputs and outputs. The mapping defined by an operator need not be "one-to-one" or "onto." If the operator determines entity  $Y$  from entity  $X$ , "one-to-one" means that there is a unique value of  $X$  for every value of  $Y$ . "Onto" means that there is some value of  $X$  corresponding to every possible value of  $Y$ .

If the mapping from  $X$  to  $Y$  is not one-to-one, certain values of  $Y$  will have several corresponding values of  $X$ . If the mapping from  $X$  to  $Y$  is not onto, certain values of  $Y$  will have no corresponding values of  $X$ . In either case, it is difficult to interpret the relationship between  $X$  and  $Y$  as one in which  $X$  is determined from  $Y$ . In other words, the relationship between  $X$  and  $Y$  contains an implicit direction:  $X$  is the input and  $Y$  is the output.

It is possible to reverse the inputs and outputs of an operator that is not one-to-one and onto by defining a new operator with certain special conventions. For example, the equation  $Y = X^2$  is ambiguous if it is viewed as a transformation from  $X$  to  $Y$ . However, if this equation is used to define an operator with  $Y$  as the input and  $X$  as the output, it can produce more than one output or no output at all, depending on the value of  $Y$ . In order to overcome this difficulty, one could define the operator with the convention that  $X$  is always the positive square root of  $Y$  if  $Y$  is greater than or equal to zero, and  $X$  is zero if  $Y$  is less than zero. Making such conventions an explicit part of an operator's definition eliminates ambiguities such as the one in this example.

## 2. Ambiguous Combinations of Entities and Operators

Although the definition of operators precludes ambiguous or impossible transformations among entities by a single operator, combinations of operators can produce insoluble models. For example, suppose operator  $O_1$  determines the value of algebraic entity  $Y$  by adding 1 to the value of entity  $X$ , and operator  $O_2$  determines the value of  $X$  by adding 1 to the value of  $Y$ . A model containing both of these operators would never reach a solution since there are no values of  $X$  and  $Y$  compatible with both operators. There are other ways that two or more operators can be inconsistent. For instance, a model might contain two operators producing different values of the same output entity. Obviously, a model designer, or an automated modeling aid, must detect and correct inconsistencies among operators if a usable decision model is to be constructed.

## 3. Connection Rules

The definitions of entities and operators, and the manner in which they are connected to each other, can be specified separately. The topological connections among operators and entities are specified by "connection rules." These rules determine exactly which entities are the inputs and outputs of each operator.

In small models it is typical for specification of the links between operators and entities to be included in the definitions of operators. For example, the demand curve for a particular commodity can be viewed as an operator transforming one specific entity (demand) into another (price). In this case the connection rule linking price and demand via the demand curve is part of the operator definition.

While separate specification of the connections between entities and operators may be cumbersome for simple problems, it can be a powerful technique for generating large models. For a typical large model with thousands of elements, a relatively small number of general operators and entities can be defined and then operator rules can be used to link them together. A connection rule might specify that one

type of operator is used many times in a model. For example, the operator representing the demand curve mentioned above might be defined in general terms with input entities that specify the exact shape of the curve. Then a connection rule could be used to specify several demand curves for various quantities.

The distinction between connection rules and operators is similar to the distinction between the definition of a variable and its specific value. Connection rules show exactly where an operator is used in a model, and the definition of the operator shows how its outputs are derived from its inputs.

When operators are defined in general terms so that they can be combined with connection rules to specify a large model, it is typical for the operator definitions to include indices. Indices are special inputs to the operator that help define its function. For example, an operator representing demand curves might include an index specifying the shape of the curve and the commodity under consideration. Several indices can be used in a single operator definition. For example, one index might specify that the demand curve is for coal and another might specify that it represents demand in 1980. Changing the indices might let the same operator specify the demand for crude oil in 1975.

Several operators and their respective inputs and outputs can be combined to form a compound operator. For example, a simple model of a coal-burning power plant might consist of a single operator that transforms the rate of fuel consumption and the energy content of the fuel into the power produced by the plant. The operator that carries out this transformation is shown in Figure 5.2(a). However, in a more detailed model of the power plant's performance, this operator may be decomposed into several others. Figure 5.2(b) shows an expanded model in which several operators and entities are used to describe the transformation defined by the original operator.

#### 4. Graphs and Trees

When operators and entities are linked together, either individually or through the use of connection rules, the result is a

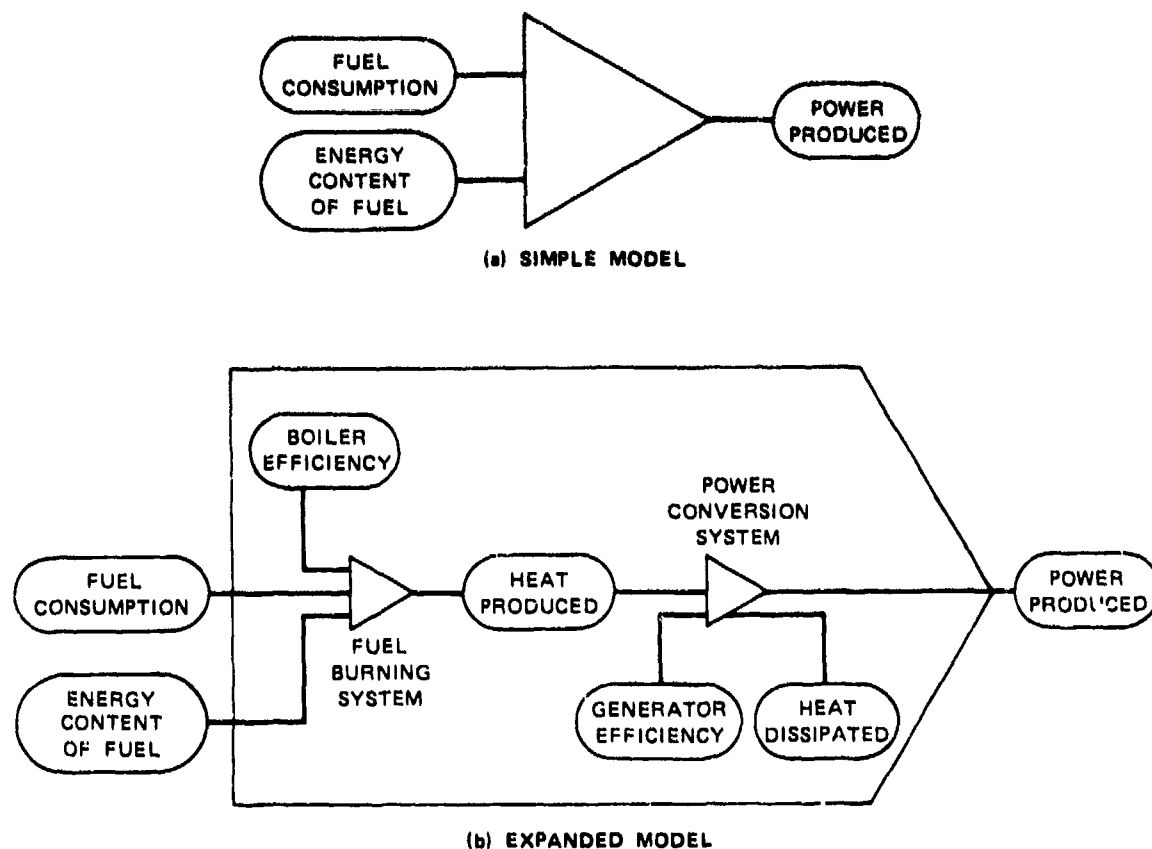


FIGURE 5.2 MODELS OF A COAL BURNING POWER PLANT

computational graph that specifies the structure of a model. Graphs can be used to specify any of the models used in decision analysis: decision trees, Markov processes, financial models, material flows over time, fault trees, etc. In fact, given the very basic nature of operators and entities, it is difficult to conceive of a model that cannot be put into the form of a graph. This is not to say that graphs are the most efficient way to model every decision problem, but rather that they are sufficiently important that they should be incorporated into any general modeling language.

Graphs can be classified according to the degree to which they are interconnected. One of the simplest forms of graphs is a tree. All the linkages of a tree are directed either from or to a unique starting point or origin. There must be only one path that connects any pair of



nodes, and if one of those nodes is the origin, all of the operators along that path must be oriented in the same direction. Figure 5.3(a) shows an example of the tree.

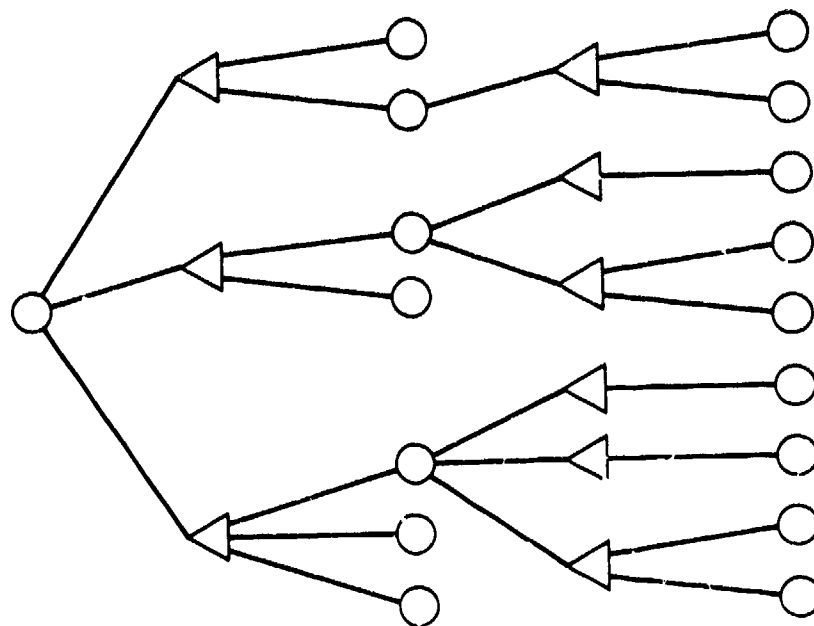
A slightly more general version of a computational graph is a coalesced tree. A coalesced tree is one in which the branches are allowed to coalesce or connect together. In a coalesced tree there can be multiple paths between each pair of nodes, but if one of the nodes is the origin, then all of the operators along all of the paths must be oriented in the same direction--either toward the origin or away from the origin. Figure 5.3(b) shows an example of a coalesced tree.

Trees and coalesced trees have the advantage of not containing loops. A graph contains a loop if two nodes in the graph are connected by two or more paths, such that all of the operators along one path lead from the first node to the second and all of the operators along another path lead in the other direction. An example of a graph containing a loop is shown in Figure 5.3(c).

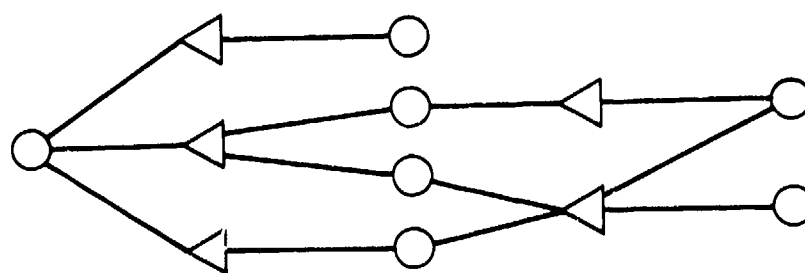
The existence of loops makes it much more difficult to process the computational graph, but it also makes it possible to simplify the structure of the graph. Conventions must be established for dealing with loops in graphs, but once this is done, it may be possible to represent very large (in fact, infinite) decision problems in compact form. For example, a Markov process can be represented as an infinitely large decision tree. Although a decision tree may be easier to process than a Markov diagram, the unbounded nature of the decision tree makes it impossible to solve completely. On the other hand, Markov processes can be represented and solved using relatively simple and compact graphs containing loops.

#### B. Function Graphs


Specific types of entities and operators can be used as applied modeling tools. One such tool is a "function graph", a graph composed of entities that are algebraic or logical variables (including arrays and vectors), and operators expressed in the form of equations. Typically, function graphs are used to construct deterministic models, but they may be extended easily to include probabilistic relationships. The following



(a) EXAMPLE OF A TREE



(b) EXAMPLE OF A LATTICE



(c) EXAMPLE OF A GRAPH WITH A LOOP

FIGURE 5.3 COMPUTATIONAL GRAPHS

discussion shows how function graphs can serve as a basis for developing decision models and how they could be implemented as part of a computerized system for decision analysis.

1. Using a Function Graph to Represent a Model for Choosing Between Alternative Missile Systems

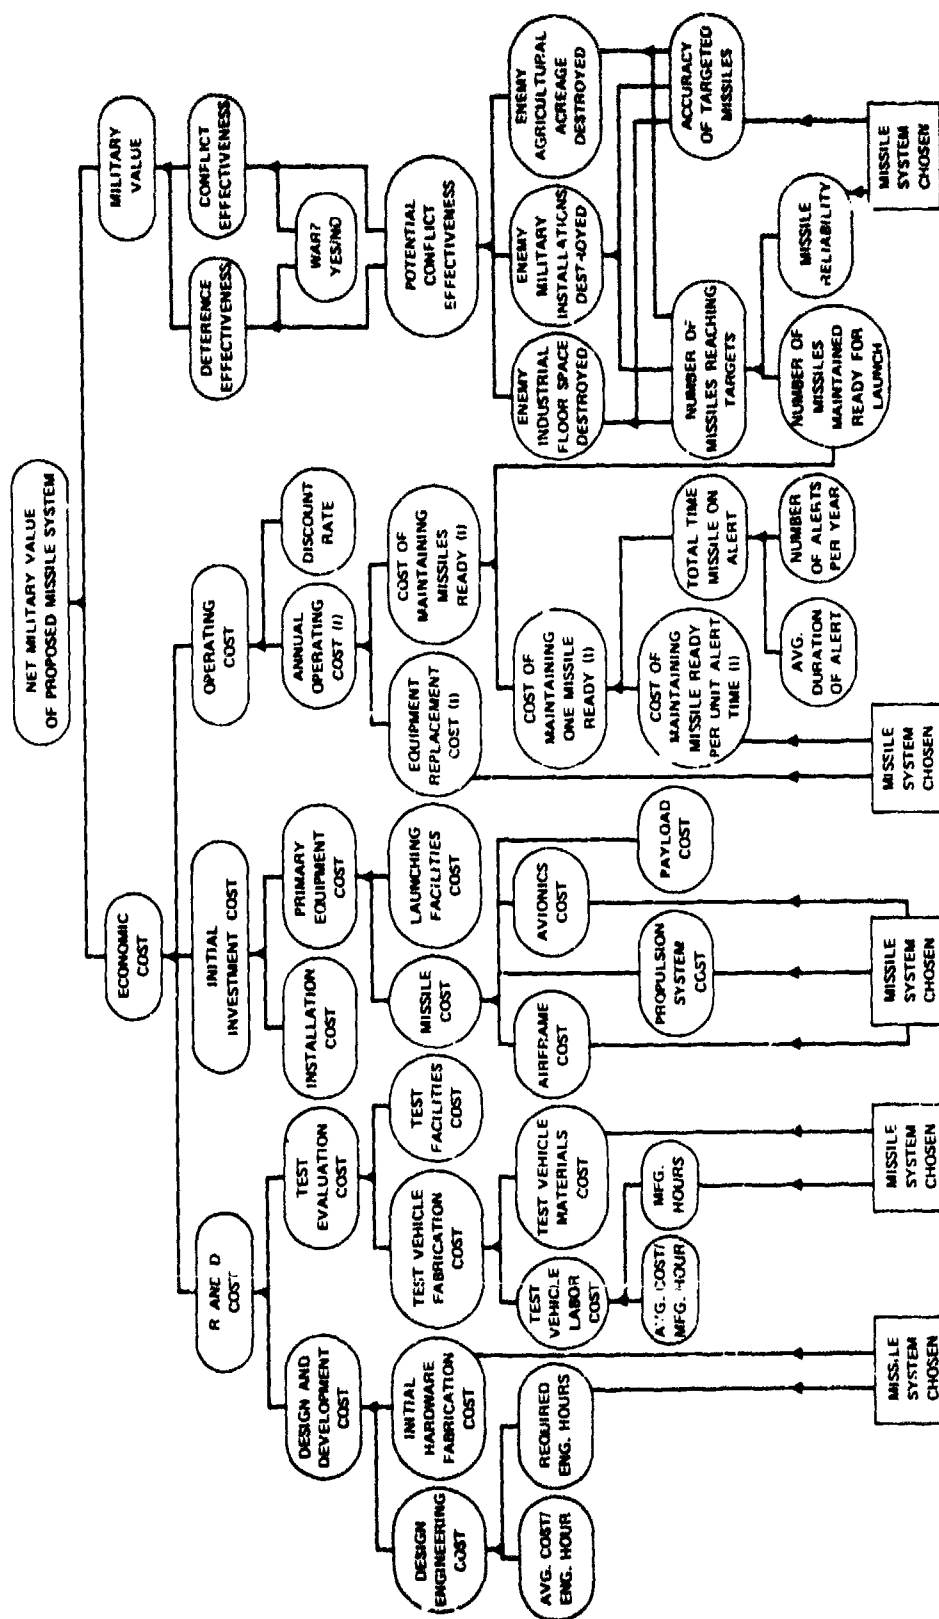
Analysis of a choice between competing weapons technology would likely result in the development of a decision model. The model would consist of important factors influencing the decision and logical and possibly algebraic dependencies among these factors. If entities are used to represent the factors represented in the decision model and operators are used to represent logical and algebraic dependencies, the decision model may be represented graphically. The proposed technique of graphical representation has been applied to a hypothetical model for choosing between new generation missile systems. A portion\* of the resulting function graph is illustrated in Figure 5.4.

In function graphs, entities are represented graphically by blocks. The name or description of the entity is written within the block. Entities nearer the top of the graph may be characterized as being more fundamentally important for judging the success or failure of a decision strategy. The entity at the top of the graph is the fundamental factor that determines the value of the decision outcome. For the model represented by Figure 5.4, the fundamental factor for judging the decision outcome is net military value, this quantity being defined in terms of a tradeoff between the military value of the missile and its economic cost. Entities nearer the bottom of the graph are the more elemental properties of the system. Those blocks which have no entering arrows are the most elemental system entities.

Triangles represent operators. A line connecting an operator to an entity means that the entity is an input or output of that operator. Thus, for example, we see from the branching structure of Figure 5.4 that initial investment cost for a missile system depends on system installation

---

\*This example will be discussed further in the next section.



An algorithm for constructing a function graph such as that of Figure 5.4 is in the next section, where we describe a computerized system designed to aid a decision maker in model development.

## 2. An Illustration of the Use of Function Graphs in a Computerized Decision Aiding System

What follows is a tentative description of the use of function graphs in an interactive computer graphics system for decision making. We shall call this overall computerized system CADMUS,\* an acronym for Computerized Aids for Decision Makers: User System. Although the characteristics we ascribe to CADMUS are within the capabilities of present generation hardware and software technology. Our objective is not to define the detailed characteristics of such a system. Instead we will describe how such a system might usefully employ function graphs as a language for model building.

To illustrate our perception of how CADMUS might use function graphs, we shall use the sample problem introduced above: choosing between alternative "next generation" strategic missiles. The problem described is purely hypothetical. It is not based on any actual analysis, nor is it meant to illustrate how such an analysis should be carried out. A real analysis of the problem would have to evaluate many factors other than those considered here.

For the purpose of illustration we assume the following hypothetical situation. Two proposals have been submitted to satisfy military requirements for a new missile design. The two missiles are the ARES and the JUGGERNAUT. While each of the proposed missiles satisfies minimum requirements, they differ somewhat in various characteristics. Estimates indicate that the JUGGERNAUT will be considerably cheaper to produce, will have a more accurate but perhaps less reliable guidance system, and will require complicated propellant servicing equipment that is more expensive

---

\*According to Greek mythology, Cadmus was a Phoenician prince reputed to have killed a dragon. Sowing the teeth of the dragon, Cadmus produced armed men that fought together until only one remained.

cost and on primary equipment cost. (This is deduced from the fact that the entity block "initial investment cost" has inputs from the blocks "installation cost" and "primary equipment cost".) We also see that the number of missiles maintained ready for launch affects the cost of maintaining missile readiness and the potential number of missiles reaching their targets. (The block "number of missiles maintained ready" has outputs to "cost of maintaining missile readiness" and "number of missiles reaching targets".)

The relationships expressed by the operators and the lines connecting them to entities are deterministic dependencies. The arrowheads formed by the orientation of the triangles on these lines define a direction for computation. The rule satisfied by the connections between operators and entities is that given fixed values for all entities with arrows leading to an operator (input entities), a unique value is specified for all entities lying along arrows leading from that operator (output entities). Thus, given installation cost and primary equipment cost, for example, the model represented by Figure 5.4 would allow initial investment cost to be calculated.

Elemental entities (those entities in Figure 5.4 with no entering arrows) may be classified as decision variables or as state variables. Decision variables are entities whose values represent decision alternatives and hence are set by the decision maker. In Figure 5.4 the elemental entities "missile system chosen" and "number of missiles maintained ready for launch" are decision variables. Decision variables are represented graphically by rectangular rather than circular blocks. The remaining elemental entities represent either parameters or state variables--entities whose values are determined by nature.

Thus, the function graph, when combined with the computational rule assignments for each operator, graphically represents the decision model's deterministic structure. Given specific decision alternatives for each decision variable and specific values for all state variables, the function graph allows the analyst to calculate the value of the fundamental entity whose value represents the net worth of the decision outcome.

to operate than similar equipment for the ARES. Initial estimates therefore indicate that, while the ARES is more expensive to produce than the JUGGERNAUT, it is cheaper to maintain at operational readiness.

To simplify the discussion we shall assume for the first level of analysis that the proposed missiles have been judged to have equal military value. This will allow us to recommend a decision between the two missiles based solely on cost, enabling us to avoid the complexities of developing model components for analyzing the military value of missile systems. As we shall discuss, however, CADMUS could be applied to the development of military value models as well as the development of cost models.

a. Using CADMUS to Build a Model

Let us imagine that CADMUS, our computerized system of decision aids, has been built. How might it be used by a decision maker interested in modeling and analyzing the missile decision problem? The decision maker begins by sitting down at a computer terminal with video screen and throwing a switch on the console indicating his request for CADMUS. A list of computer aids for decision analysis appears instantly on the screen, and our decision maker indicates his wish for assistance in deterministic model development by touching an appropriate item in this list with the light pen attached to his terminal.

Instantly, the video screen is transformed into the control panel illustrated in Figure 5.5. The center of the control panel, which displays messages from the computer system, may also be used as a "scratch pad" for drawing, typing, or writing. As the light pen is pulled across this area, a displayed "ink" track appears to flow from the pen. Items appearing here may be erased by scrubbing over them with the pen. Around the edges of the screen are various control "pushbuttons." If one of these is "pushed" (by touching the pen to it), the system performs the indicated action. The buttons in the lower right hand corner of the screen form a "calculator" for writing mathematical or logical equations. The buttons on the left of the screen are for system control.

EXPAND	CONTRACT	ROLL UP	ROLL DOWN	ROLL LEFT	ROLL RIGHT	ERASE	HARD COPY	STORE PAGE	RECALL PAGE	DATA	MODEL	PLOT	SENSITIVITY	CHANGE RATE	PROB DIST																
<div>DEPENDENT VARIABLE</div> <div>INDEPENDENT VARIABLE</div> <div>FROM</div> <div>TO</div> <div>STEP SIZE</div>																															
<div>=</div> <div>—</div> <div>+</div> <div>×</div> <div>÷</div> <div>1</div> <div>LOG</div> <div>LN</div> <div>7</div> <div>8</div> <div>9</div> <div>SIN</div> <div>COS</div> <div>TAN</div> <div>4</div> <div>5</div> <div>6</div> <div>1/x</div> <div>1</div> <div>2</div> <div>3</div> <div>0</div> <div>.</div> <div>AND</div> <div>IF</div> <div>OR</div>																															

FIGURE 5.5 CONTROL PANEL USED FOR MODEL FORMULATION BY CADMUS



The following question appears at the center of the video screen:

WHAT QUANTITY OR QUANTITIES WOULD YOU NEED TO KNOW TO EVALUATE THE OUTCOME OF YOUR DECISION?

After thinking a moment, our decision maker decides that he would need to know two things to accurately evaluate the missiles. First, he would need to know the total economic cost incurred in the construction and operation of each missile system. Second, since military decisions cannot be made on the basis of cost alone, he would have to know the total military value to be derived from each missile system. He therefore types ECONOMIC COST and then MILITARY VALUE on the computer terminal; as he does, the words appear on the video screen. Economic costs will be fairly easy to define; military value will be considerably harder. As stated above, we shall limit ourselves to an analysis of economic cost. Following the development of an economic cost model we shall describe briefly how a military value model would be constructed.

The decision maker indicates that he wishes to explore economic costs by touching the words ECONOMIC COST with the light pen. CADMUS responds by flashing the following question on the screen.

SUPPOSE YOU HAD A CRYSTAL BALL. YOU MAY ASK IT ANY NUMERICAL QUESTION EXCEPT "WHAT IS 'ECONOMIC COST'?" FOR WHAT QUANTITY OR QUANTITIES WOULD YOU ASK IN ORDER TO CALCULATE "ECONOMIC COST"?

Our decision maker knows that there are three major cost categories for weapon systems: research and development costs, initial investment costs, and operating costs. He types R&D COST, INITIAL INVESTMENT COST, and OPERATING COST. Immediately, the video screen represents these entries as entities in a graphical structure as shown in Figure 5.6.

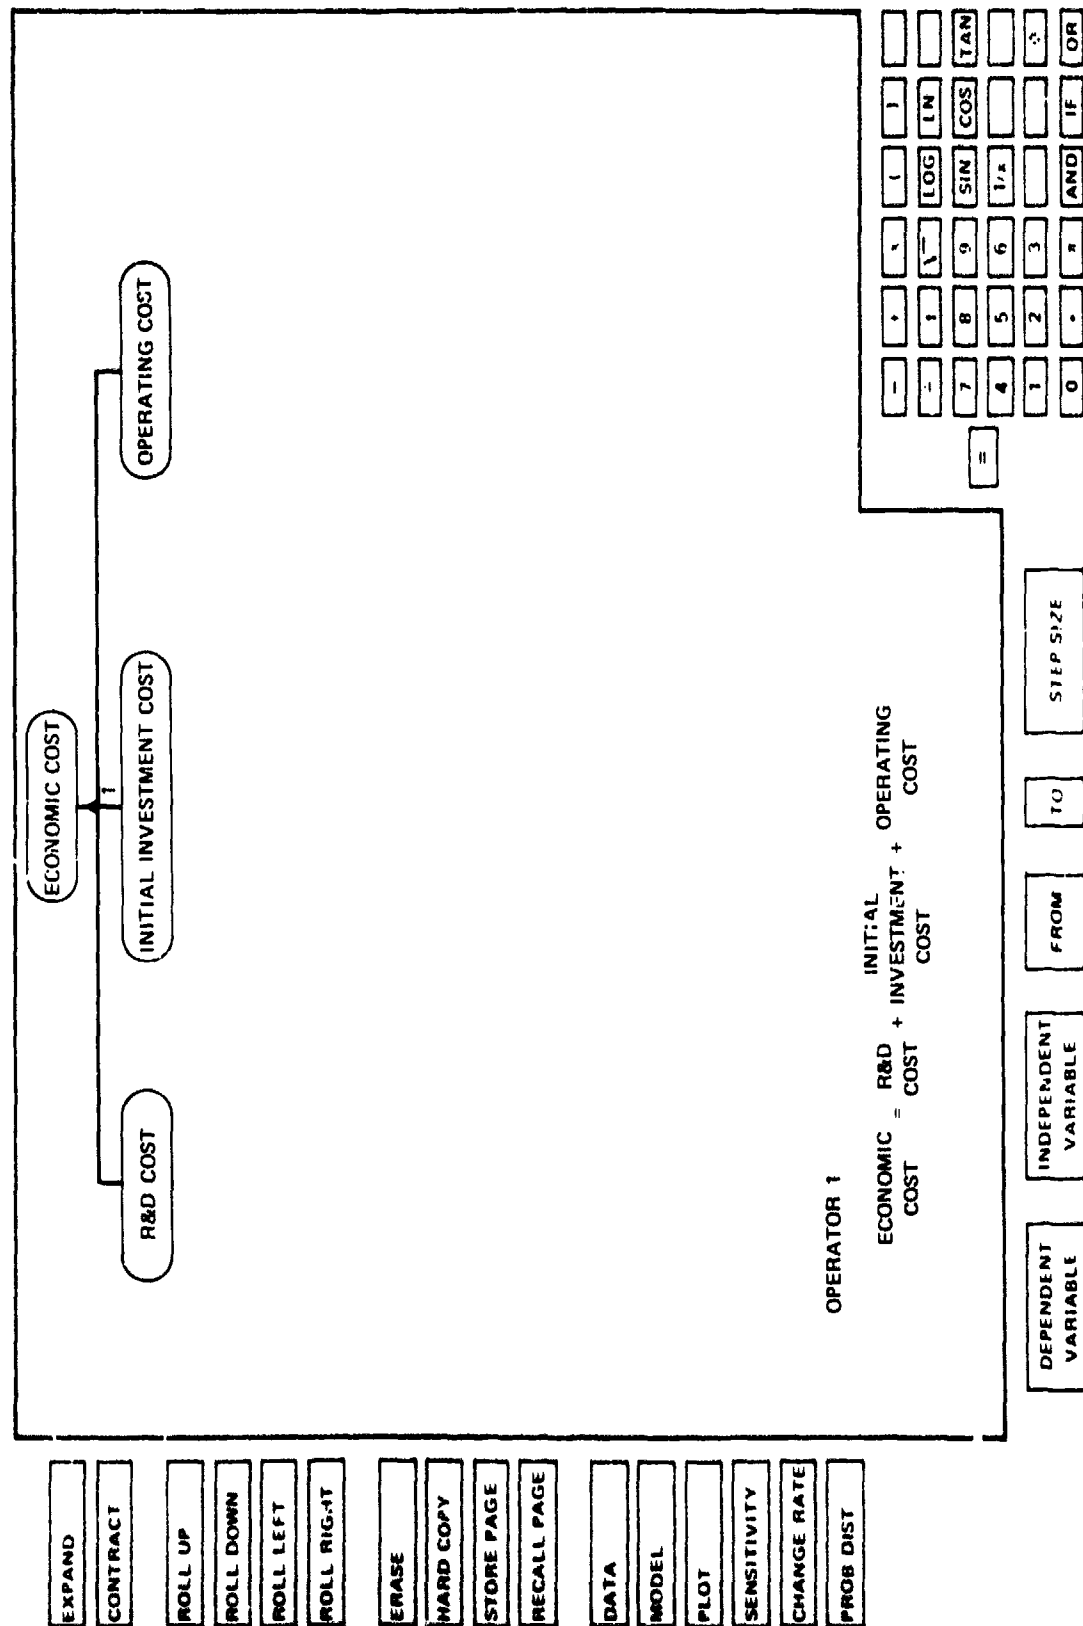


FIGURE 5.6 FIRST STEP FOR MODELING THE MISSILE DECISION

The relationship between the three entities is represented graphically by an "operator" triangle to which CADMUS has assigned the number 1. To specify an analytic definition for the operator the decision maker turns to the calculator in the lower right hand corner of the screen. Thinking "economic cost will be the sum of R&D, investment, and operating costs," he uses his light pen to write\*

$$\begin{array}{c} \text{TOTAL} \\ \text{ECONOMIC} \\ \text{COST} \end{array} = \begin{array}{c} \text{R\&D} \\ \text{COST} \end{array} + \begin{array}{c} \text{INITIAL INVEST-} \\ \text{MENT COST} \end{array} + \begin{array}{c} \text{OPERATING} \\ \text{COST} \end{array}$$

by alternately touching the appropriate entity blocks in the branching structure and operator buttons on the calculator. As he proceeds, the equation defining operator number one appears at the bottom of the video screen.

The decision maker touches the light pen to the entity block marked "operating cost" and CADMUS answers:

SUPPOSE YOU HAD A CRYSTAL BALL. YOU MAY ASK ANY NUMERICAL QUESTION EXCEPT "WHAT IS 'OPERATING COST'?" FOR WHAT QUANTITY OR QUANTITIES WOULD YOU ASK IN ORDER TO CALCULATE "OPERATING COST"?

"Well," thinks our decision maker, "for the basis of comparison we have specified that the missile maintain a state of operational readiness for a four-year period." He decides to express operating costs on an annual basis and to provide for the possibility of using a discount rate to discount the magnitude of future operating expenditures. If  $r$  is the discount rate, the present value of four years of operating costs is given by :

---

\*As mentioned above, a user may "push" a "button" illustrated on the video screen by touching that button with the light pen. In the text we represent the pushing of a button by enclosing the button label within a rectangular figure. Permanent control buttons are shown as rectangular figures with sharp corners. User-defined entity buttons are enclosed in figures with rounded corners. Rectangular blocks are not placed around those items that are entered through the terminal keyboard.

$$PV = rOC_1 + r^2OC_2 + r^3OC_3 + r^4OC_4$$

where  $OC_1$  is the operating cost in year 1. Notice that if  $r$  is set to one, operating cost is obtained by simply adding the annual operating costs for each of the four years of operation. If  $r$  is less than one, then operating costs in future years are not counted so heavily as current expenditures. To input the above equation the decision maker finds it convenient to define the entity annual operating costs as a vector whose components are the "annual operating costs" in each of the next four years. He types DISCOUNT RATE and then ANNUAL OPERATING COSTS(I). The "I" will be used to denote the year to which the annual operating cost corresponds. The first year will be denoted by  $I = 1$ , the second by  $I = 2$ , and so forth. The branching structure momentarily disappears and then reappears as shown in Figure 5.7.

Using the light pen, the entity blocks shown in the function graph and the "calculator" in the left-hand corner of the screen, our decision maker "pushes" the buttons that express operating cost as the present value of the four annual operating costs:\*

$$\begin{aligned} \text{OPERATING COST} &= \text{ANNUAL OPERATING COST (1)} \times \text{DISCOUNT RATE} \\ &+ \text{ANNUAL OPERATING COST (2)} \times \text{DISCOUNT RATE}^2 \\ &+ \text{ANNUAL OPERATING COST (3)} \times \text{DISCOUNT RATE}^3 \\ &+ \text{ANNUAL OPERATING COST (4)} \times \text{DISCOUNT RATE}^4 \end{aligned}$$

The equation appears near the bottom of the screen as shown in Figure 5.7.

\*The up arrow button  $\uparrow$  denotes that the quantity immediately following is to be interpreted as an exponent.

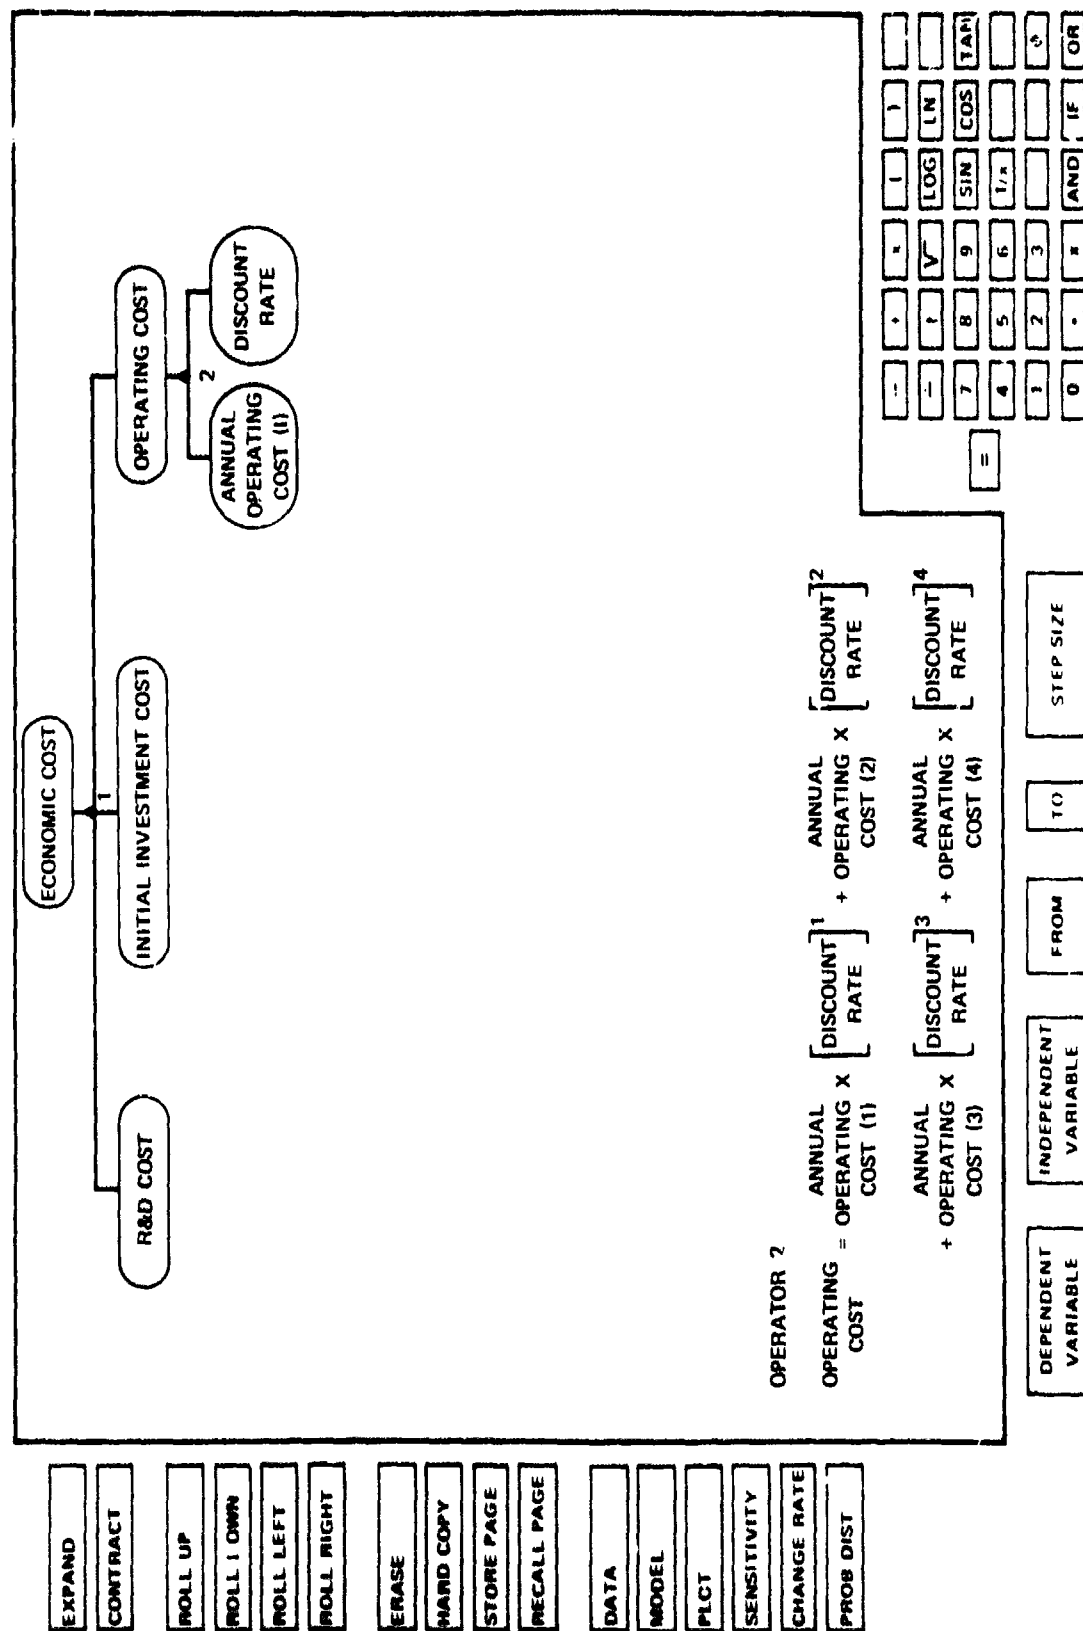


FIGURE 5.7 SECOND STEP FOR MODELING THE MISSILE DECISION

Next our decision maker touches his light pen to the block marked "annual operating cost (I)." The familiar question, appears on the screen this time asking him what he would need to know to define the annual operating cost in year I. The decision maker types EQUIPMENT REPLACEMENT COST (I) and COST OF MAINTAINING MISSILES READY (I), thinking these to be the two main components, and uses the calculator to define annual operating costs in year I as the sum of these quantities. The video screen now appears as in Figure 5.8. Similarly, he defines the cost of maintaining missile readiness as the product of the cost of maintaining one missile ready times the number of missiles maintained ready. The screen appears as in Figure 5.9.

Notice that while the decision maker began by thinking about very aggregated concepts such as total military cost of the proposed missile, the branching technique leads him to consider increasingly specific items, such as the cost of maintaining a single missile ready and the number of missiles maintained ready, which are much easier to estimate. Our decision maker may continue to use the branching technique until he has defined all entities in terms of elemental quantities that he feels comfortable estimating. As illustrated by the definition of operator 2 in Figure 5.7, operators need not be confined to simple arithmetic functions like addition and multiplication. If necessary more general functions or subroutines may be used to define the relationships among system entities.

#### b. The Completed Cost Model

For the purposes of our example, let us suppose that our decision maker continues expanding the system model by branching from entities until the graph appears as shown in Figure 5.10. We imagine at this point that the decision maker has decided that with a little thought he can come up with estimates of values for the elemental state variable entities, and he sees that the important decision alternatives-- "missile system chosen" and "number of missiles maintained ready"--have been represented. Of course, the structure shown in the figure is unrealistic because it is overly simplified and some important considerations have been omitted. In a real analysis the decision maker would

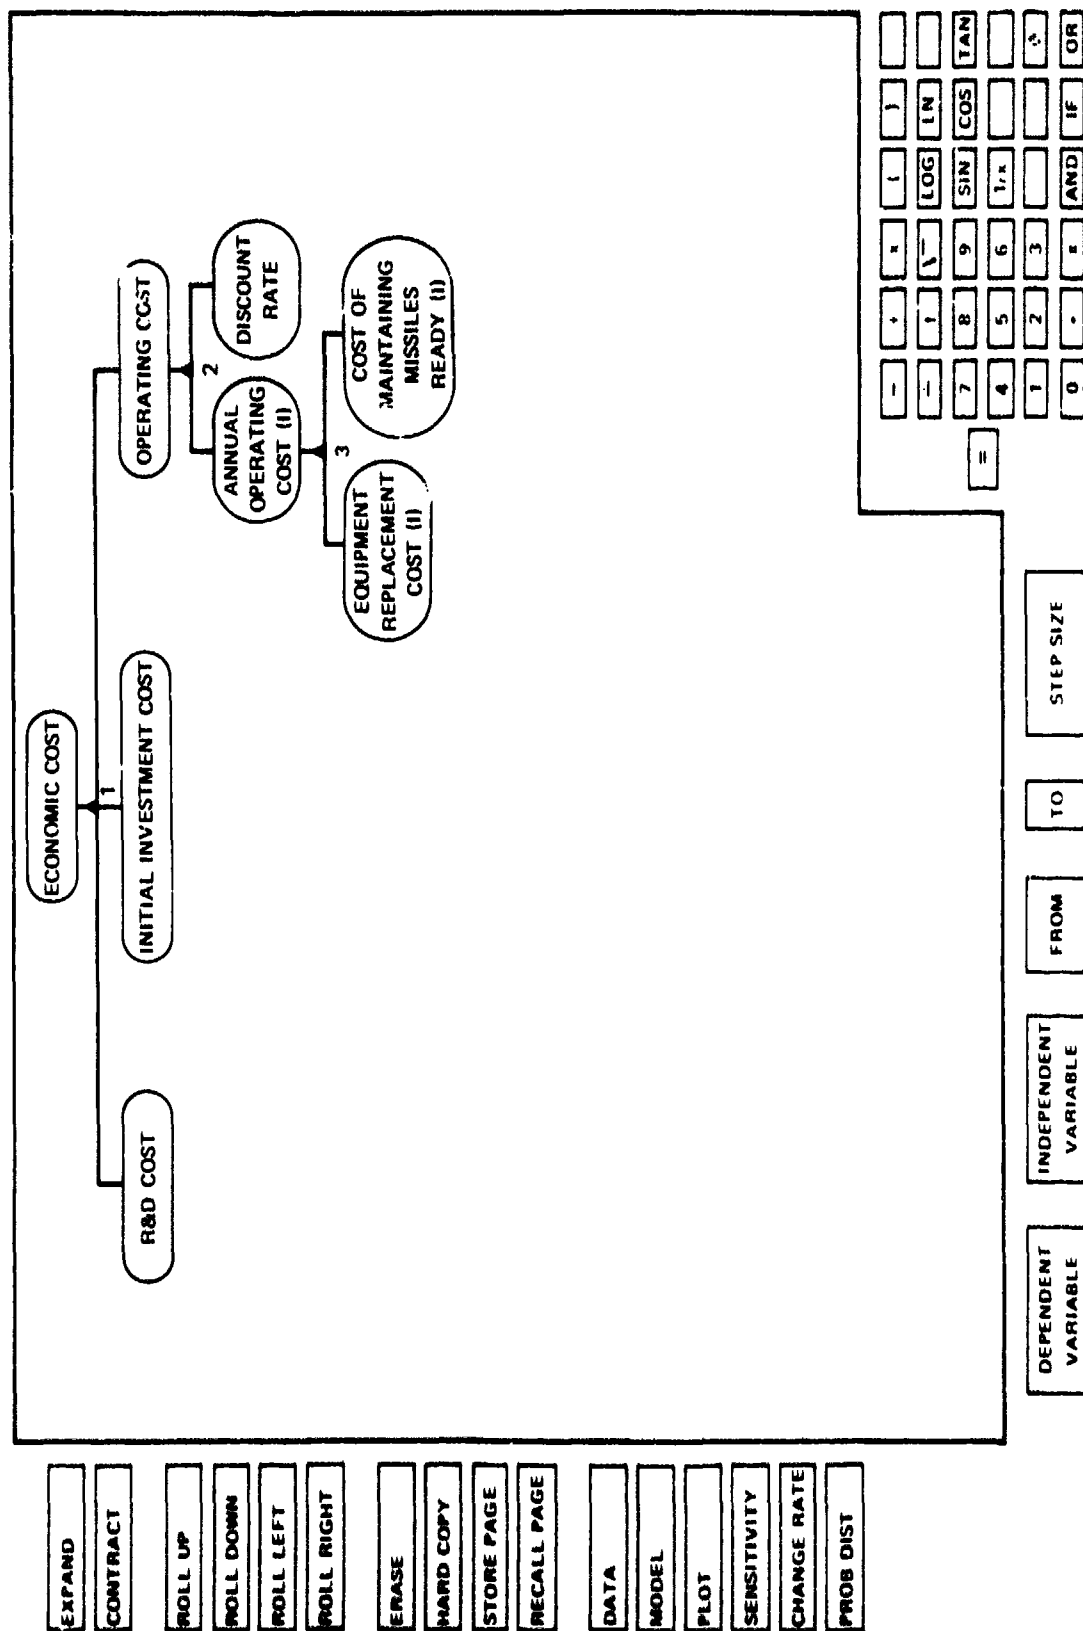


FIGURE 5.8 THIRD STEP FOR MODELING THE MISSILE DECISION

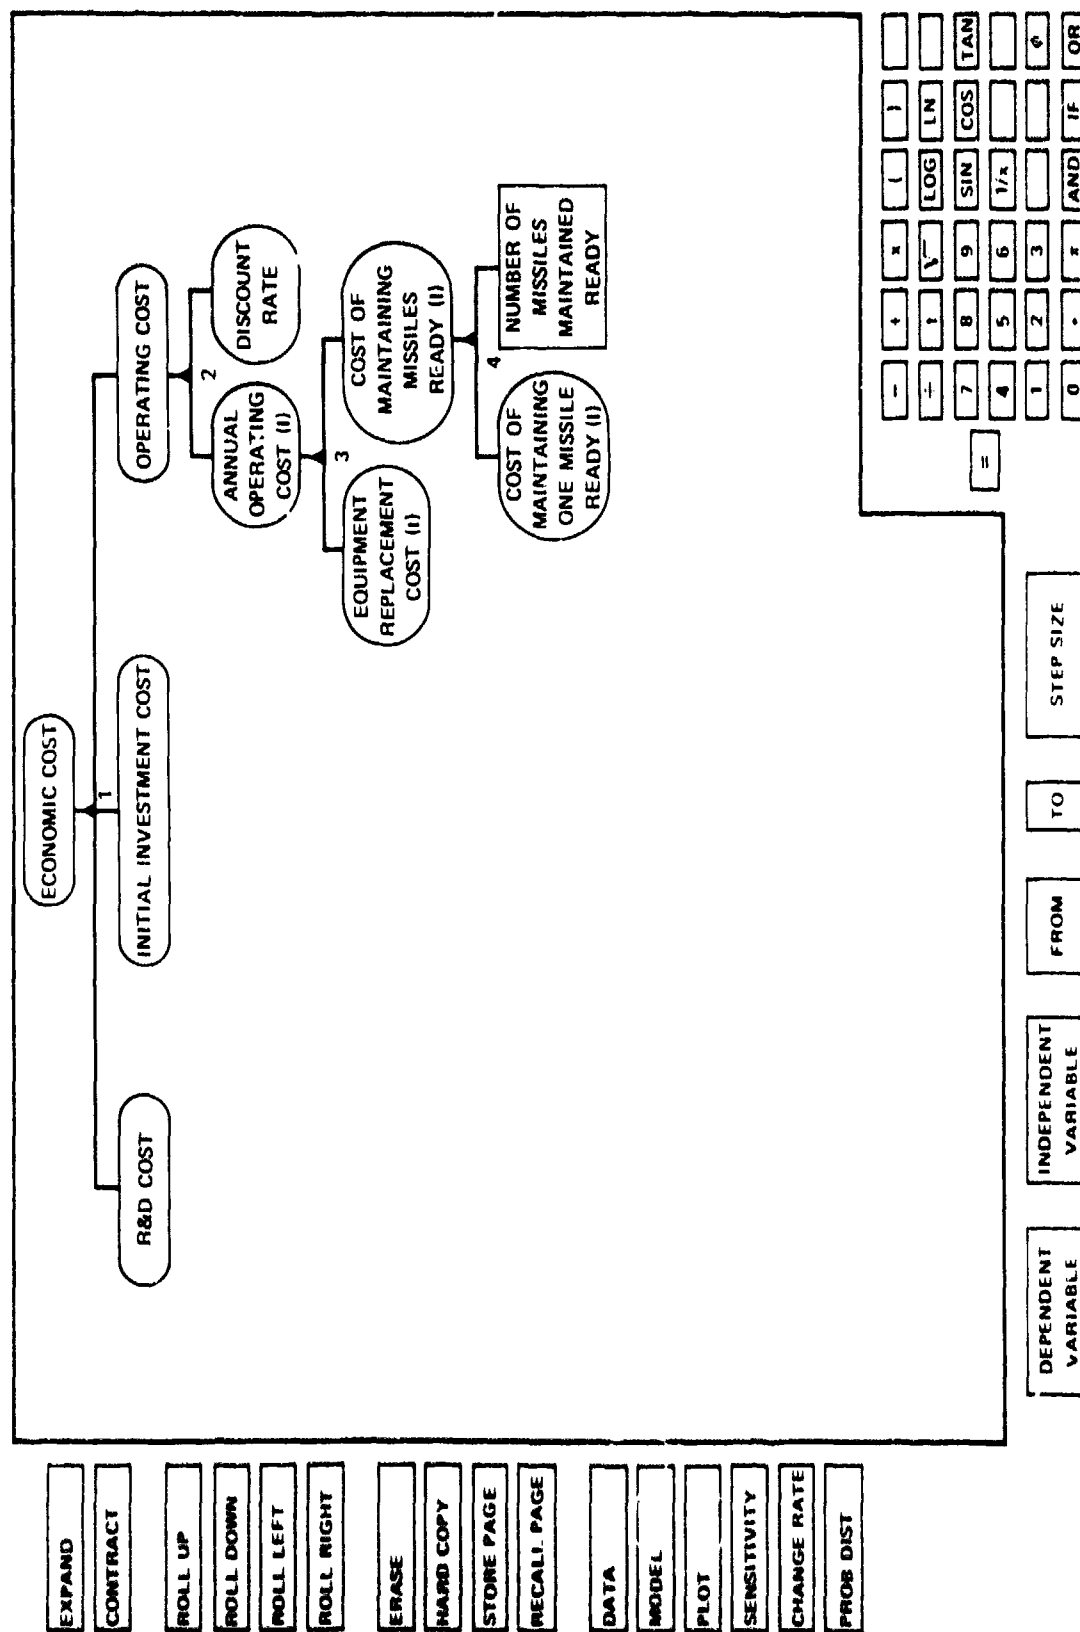


FIGURE 5.9 FOURTH STEP FOR MODELING THE MISSILE DECISION



undoubtedly wish to expand the structure to a considerably higher level of detail than that shown in Figure 5.10. Nevertheless, the function graph of the figure will suffice for illustration.

To avoid the necessity of modeling the military value of the missile systems we have assumed that the competing missiles have been judged to have equal military value. However, similar procedures to that described above could be used by the decision maker to construct a branching structure representing the military value of the missile system chosen. The decision maker may feel, for example, that the value of the proposed missile system should be assessed in two areas. First, it should provide deterrence to a thermonuclear war. Second, if war does come the new missile should help us to minimize damage to ourselves and speed a favorable military outcome of the conflict. The function graph developed to model the military value of the missile system might look something like that shown in Figure 5.4.

To specify the operator that maps the various military objectives into military value, each objective could be given a relative weight. This could be done by constructing tradeoff curves that illustrate the decision maker's willingness to trade one objective off against another. The operators in the function graph representing military value would thus be defined by mathematical relationships that are primarily subjective in nature. For this reason they might be more difficult to determine than those used to determine economic cost, but the principle is the same.

When combined with analytic operator definitions the function graph shown in Figure 5.10 completely specifies the deterministic structure necessary for evaluating the economic cost of various decision strategies. Given specific alternatives for the decision variables "missile system chosen" and "number of missiles maintained ready," and given specific values for the elemental state variables, evaluation of the function graph will produce a total economic cost.

The function graph, therefore, represents a mathematical cost model for the missile decision. An important question is, "How good is this model?" CADMUS supplies a number of tests for model evaluation.

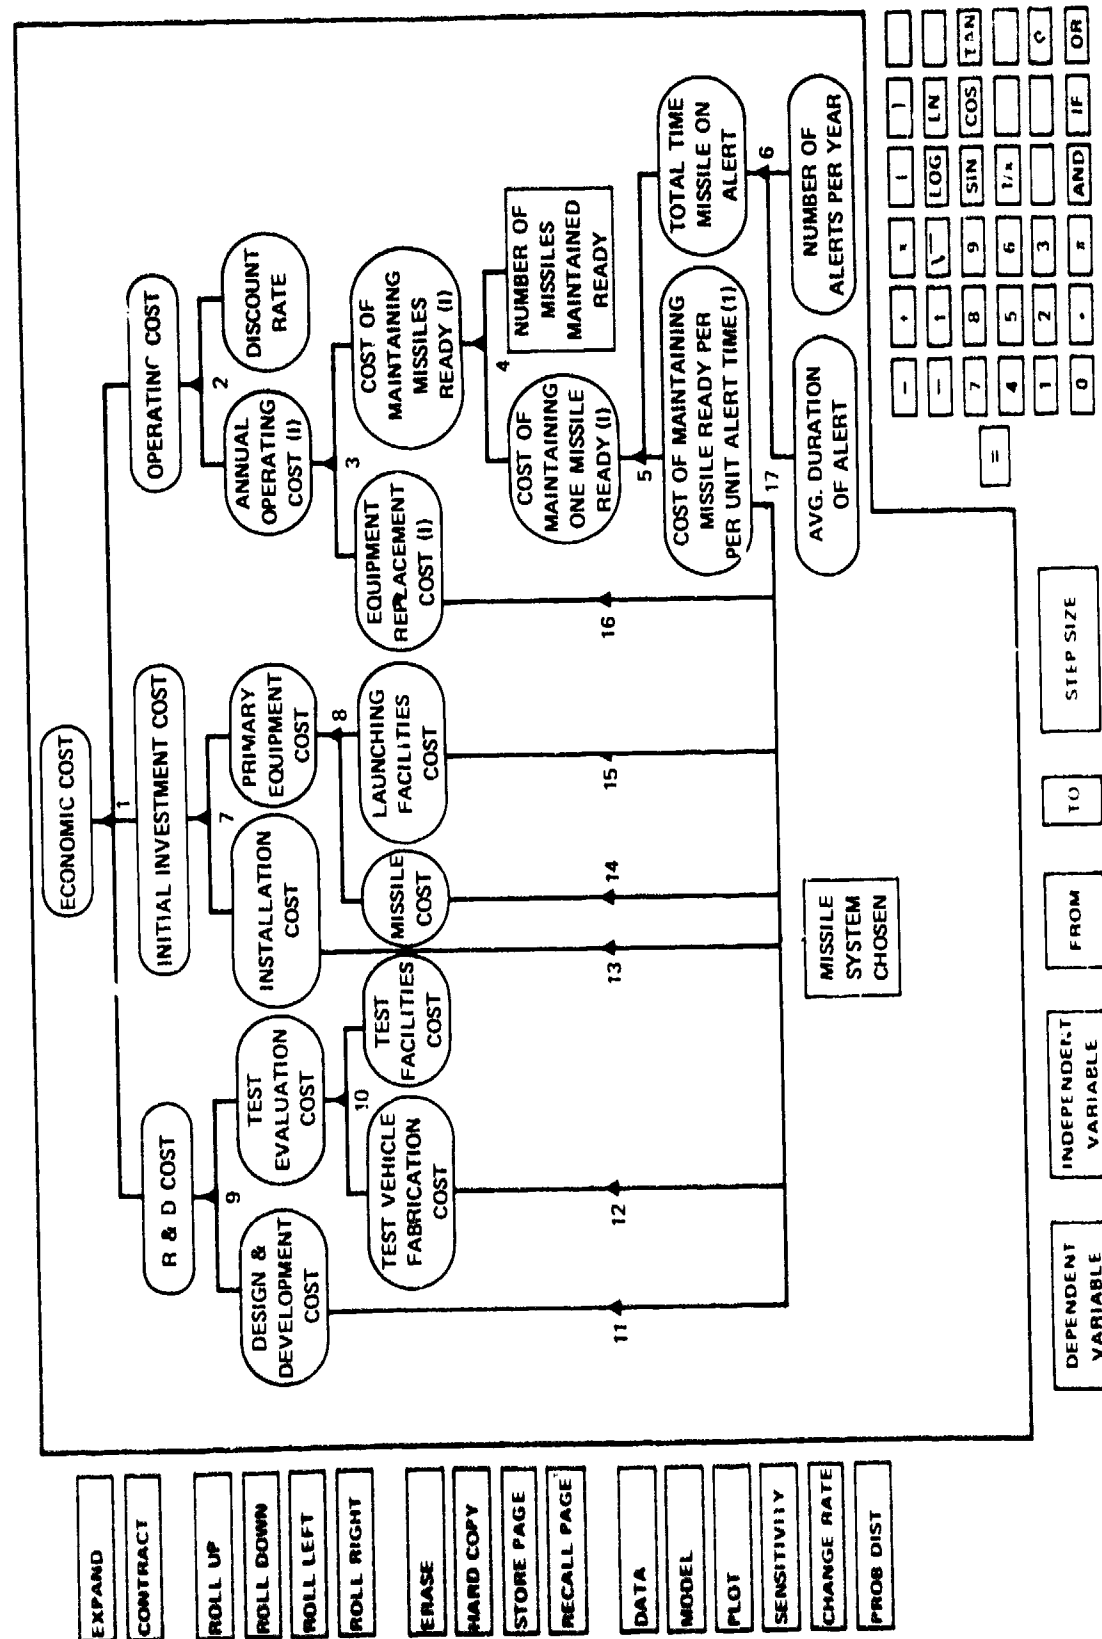


FIGURE 5.10 FINAL STEP FOR MODELING THE MISSILE DECISION

c. Model Analysis

CADMUS has the ability to perform sensitivity analyses. In sensitivity analysis, the decision analyst tries to determine the change in the model's selection of alternative actions or outcome values that would result from a given change in the model's assumptions. Assumptions that produce small changes are apparently relatively insignificant, while assumptions that produce considerable changes are likely quite significant. CADMUS provides for a number of automatic sensitivity calculations. For example, suppose our decision maker felt uneasy about his estimate of the number of times the missile will be placed on alert and therefore he is uncertain about the amount of time the missile must maintain a state of readiness. Naturally, the total time the missile must maintain a state of readiness strongly influences missile system operating costs. He might feel that the number of alerts could add up to anything between zero and fifty per year. Knowing how much this uncertainty contributes to uncertainty about total economic cost will be important in determining whether further effort should be expended to clarify the prediction of future missile alerts.

Suppose therefore that our decision maker wants to calculate how total economic cost changes as the number of alerts is varied. He begins by choosing a specific alternative for the decision variable:

(MISSILE SYSTEM CHOSEN) ☐ JUGGERNAUT

Control buttons for the two alternatives, ARES and JUGGERNAUT, will have been defined when the operators 11 through 17 in Figure 5.10 are defined. Next nominal values are assigned to elemental state variables:

(AVG. DURATION OF ALERT) ☐ 20 Minutes  
(EQUIPMENT REPLACEMENT COST) ☐ 2.8 , 1.5 , 1.6 , 1.8 MILLION  
(DISCOUNT RATE) ☐ 1.1\*

\*To simplify notation, numbers that require a sequence of button pushing operations will be represented as if they could be specified by a single button. For example, the notation ☐ 1.1 means that the buttons ☐ 1 ☐ . ☐ 1 are pushed in succession.

By touching his light pen to the **SENSITIVITY** control button, the decision maker signals CADMUS that he wishes to perform a sensitivity calculation. Pushing **DEPENDENT VARIABLE** and then **ECONOMIC COST** establishes economic cost as the dependent variable in the relationship. Our decision maker wishes to learn how the value of this variable varies as he varies the number of yearly missile alerts between the values of zero and 50. To communicate this to CADMUS he uses his light pen to push the respective buttons **INDEPENDENT VARIABLE** and **NUMBER OF ALERTS PER YEAR** and then FROM **0** TO **50**. The video screen momentarily goes blank and then appears as Figure 5.11, showing the functional relationship established between the variables "total economic cost" and "number of alerts per year"\*

The decision maker may wish to test this same sensitivity assuming instead that the ARES missile is chosen. If he were to set "missile system chosen" to "ARES" and re-run the sensitivity, the result might look something like that shown in Figure 5.12. An instructive exercise would be to display simultaneously the sensitivity plots for the two decision alternatives. By storing the results and then pushing RECALL, our decision maker could generate the plot shown in Figure 5.13\*\*. Such a plot illustrates the closed loop sensitivity of "economic cost" to "number of alerts per year." The closed loop sensitivity shows how the optimal value of the dependent variable changes with different values for the independent variable. Assuming that the military value of the ARES and JUGGERNAUT are identical, the objective is to choose the decision strategy that minimizes "economic cost". The closed loop sensitivity is therefore given by the piecewise linear curve formed by the lower envelope of the two straight line curves. If the number of alerts is anticipated to be roughly greater than 27 per year, the system with the lower operating cost (ARES) should be chosen. Otherwise the JUGGERNAUT should be chosen.

\*The numerical values presented here are purely hypothetical.

\*\*Since our objective here is not to define the operation of CADMUS in detail, the mechanics of instructing the system to store and recall graphic information are not discussed.

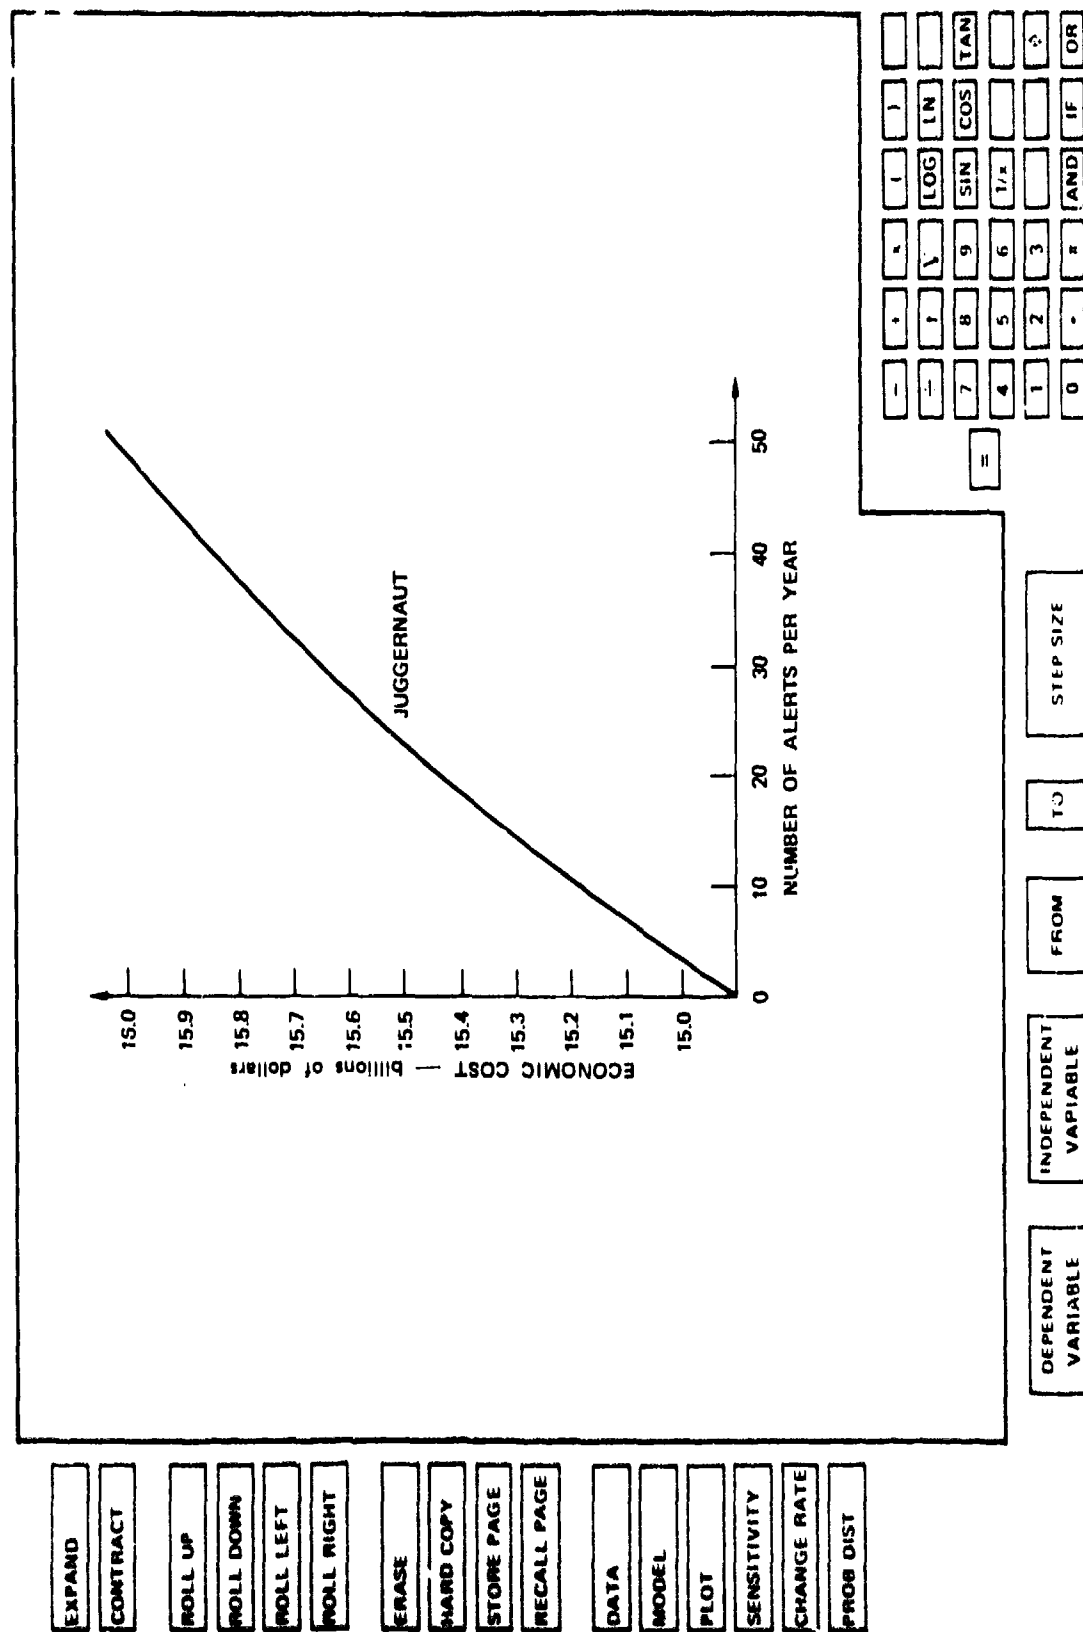


FIGURE 5.11 SENSITIVITY OF ECONOMIC COST TO NUMBER OF ALERTS PER YEAR (JUGGERNAUT SYSTEM)

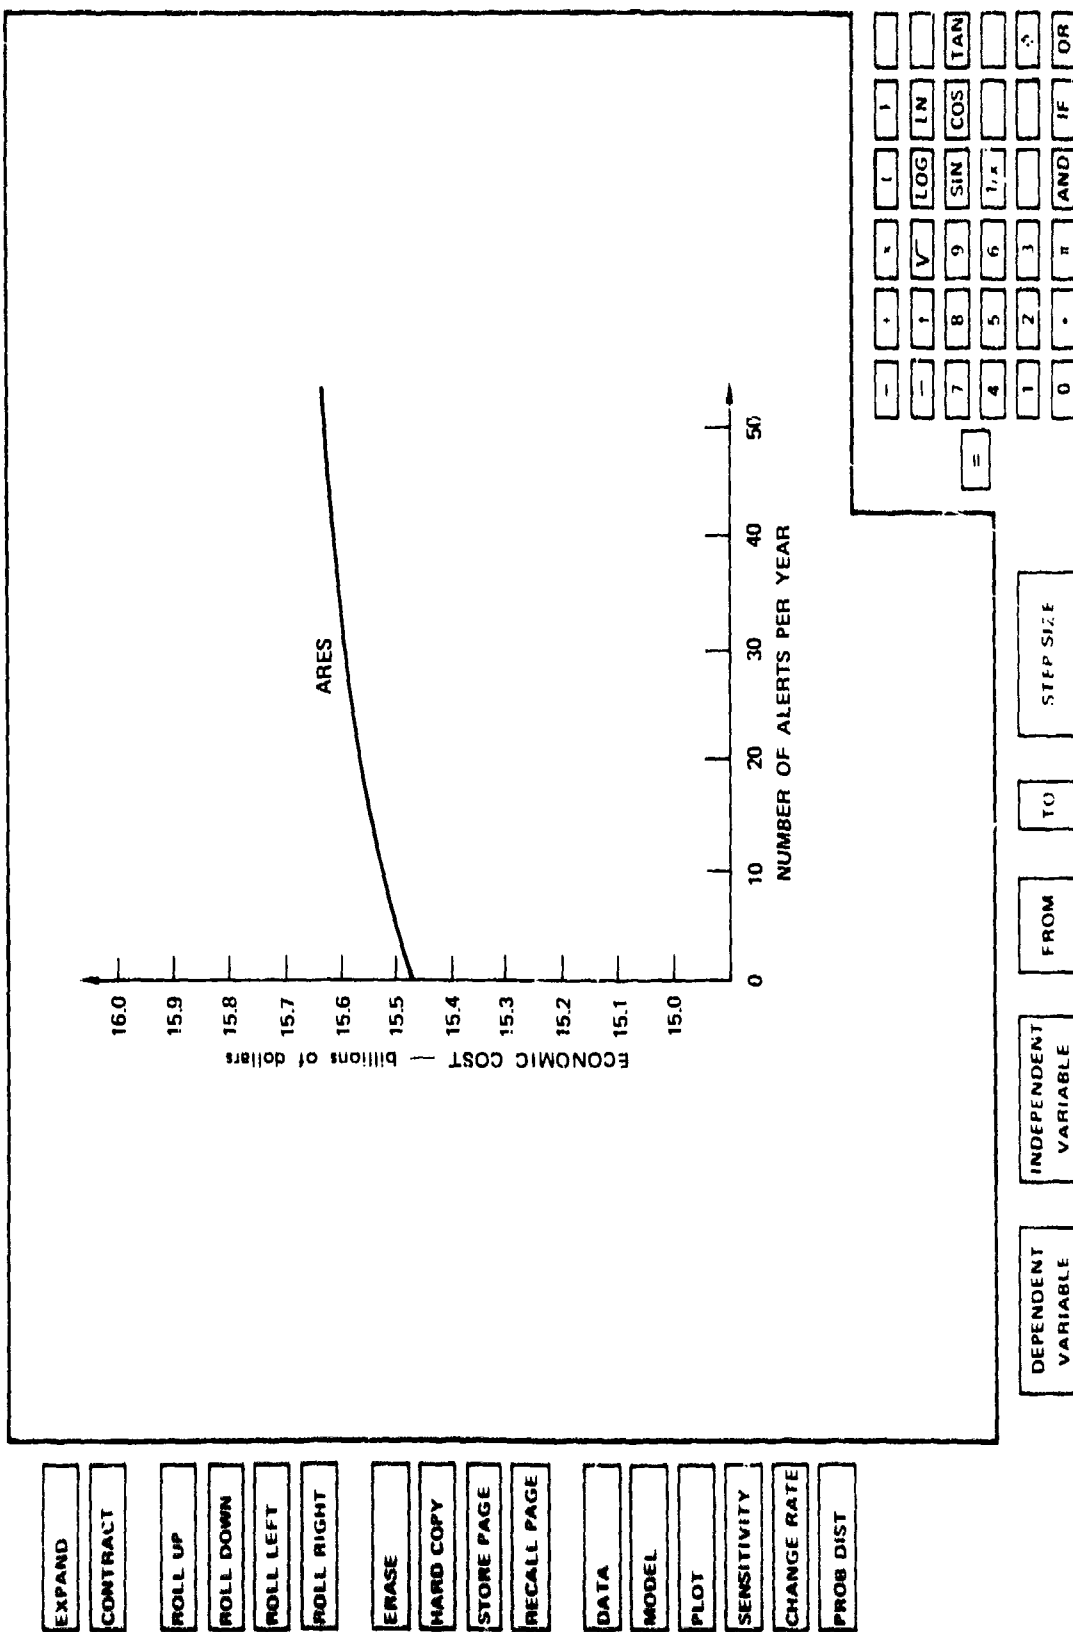


FIGURE 5.12 SENSITIVITY OF ECONOMIC COST TO NUMBER OF ALERTS PER YEAR (ARES SYSTEM)




FIGURE 5.13 SENSITIVITY OF ECONOMIC COST TO NUMBER OF ALERTS PER YEAR (JUGGERNAUT AND ARES SYSTEMS)

CADMUS may be requested to perform a number of additional calculations which will enable the decision maker to evaluate his alternative decision strategies. For example, the decision maker may ask CADMUS to check for dominated courses of action, that is, decision strategies that for all possible state outcomes yield values that are equaled or exceeded by other alternative strategies. The decision maker may ask to see the "outcome lottery" associated with a given course of action. CADMUS will then supply the cumulative probability distribution that shows the probability that the outcome associated with that decision strategy will have a value less than or equal to any given amount.

To illustrate, suppose that our decision maker were to specify a probability distribution for the number of alerts per year such as that shown in Figure 5.14. CADMUS would provide a number of methods to aid the decision maker in assessing such a distribution. He would also specify probability distributions for all other uncertain variables that strongly influence economic cost. Then, by touching his light pen to PROB DIST ("probability distribution" control button) and then to ECONOMIC COST, the decision maker would instruct CADMUS to produce cumulative probability distributions for "economic cost" under each decision alternative. The system would respond as shown in Figure 5.15. The curve associated with a given missile system shows the probability that its total economic cost would fall below any given amount.

### C. Decision Trees and Influence Diagrams

This section describes a novel approach to probabilistic modeling, one that builds on a fundamental conceptual tool of decision analysis--decision trees--and introduces a related concept--influence diagrams. Together, these conceptual tools constitute a basic foundation for a system of probabilistic modeling aids, in the same way that structural graphs form the basis for deterministic modeling. In fact, further research may show that a synthesis of these concepts will provide a general tool for all types of modeling.



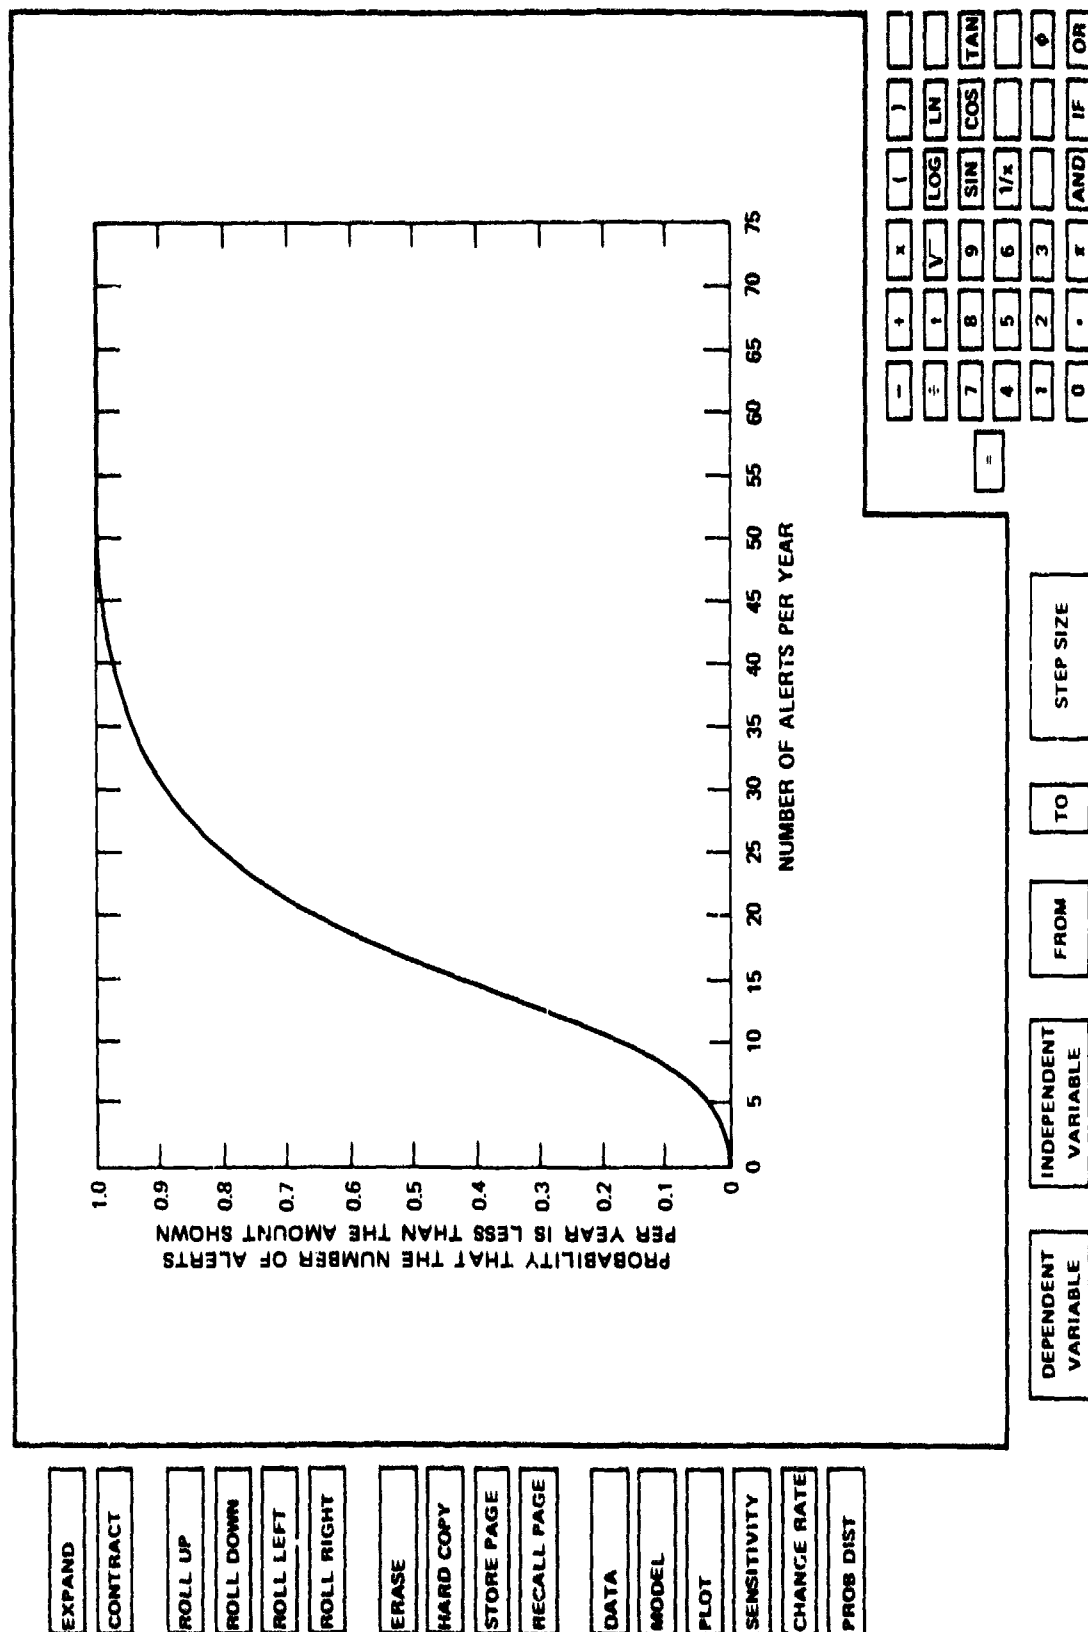


FIGURE 5.14 PROBABILITY DISTRIBUTION ILLUSTRATING UNCERTAINTY IN THE NUMBER OF ALERTS PER YEAR

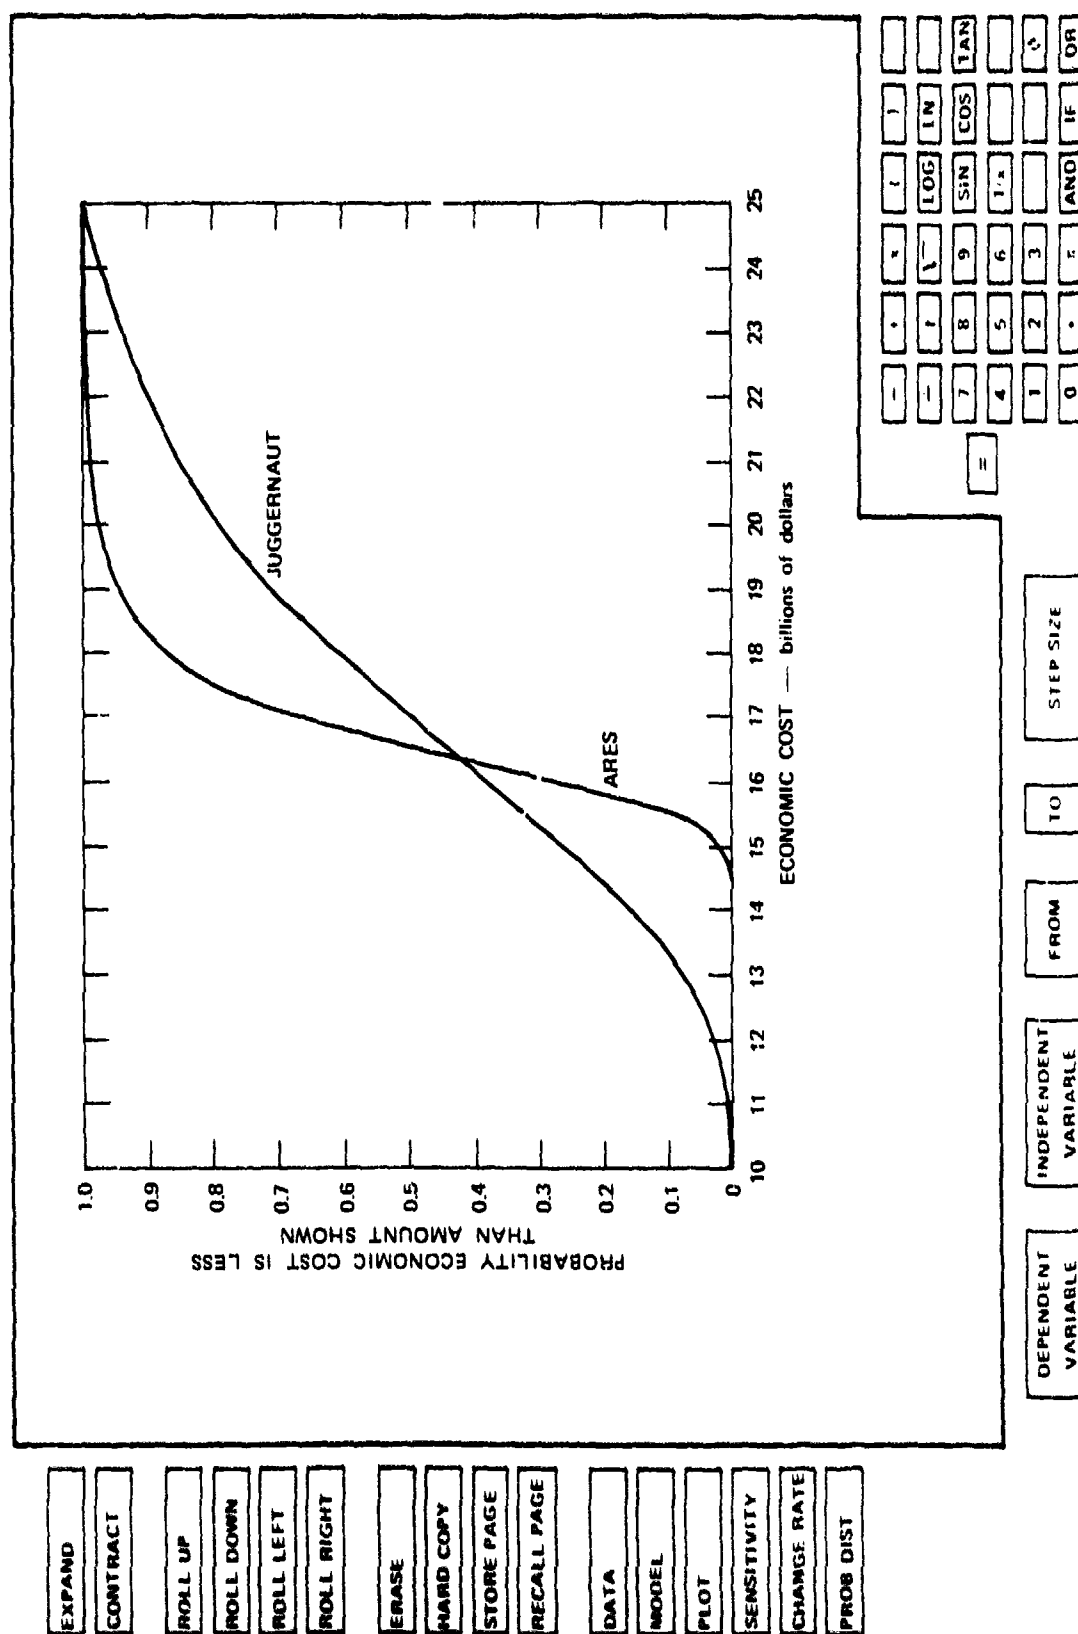


FIGURE 5.15 PROBABILITY DISTRIBUTIONS ILLUSTRATING UNCERTAINTY IN ECONOMIC COST FOR TWO PROPOSED MISSILE SYSTEMS

## 1. Properties of Decision Trees and Influence Diagrams

One of the most perplexing aspects of making decisions under uncertainty is the problem of representing and encoding probabilistic dependencies. A probabilistic dependency is one that arises as a result of uncertainty. For example, if  $a$  and  $b$  are known variables and  $c = a + b$ , then it is clear that  $c$  depends on both  $a$  and  $b$ , both in a vernacular sense and in a mathematical sense. However, suppose  $a$  is known and  $b$  is uncertain. Then  $c$  is probabilistically dependent on  $b$  but not on  $a$ . The reason is that knowing the specific value of  $b$  tells us something new about  $c$ , but there is no such possibility with respect to  $a$ .

### a. Probabilistic Independence

Probabilistic independence, like the assigning of probability itself, depends on the state of information possessed by the assessor. Let  $x$ ,  $y$ , and  $z$  be aleatory state variables of interest, which can be either continuous or discrete. Then  $\{x|S\}$  is the distribution assigned to  $x$  given the state of information  $S$ . Two variables  $x$  and  $y$  are probabilistically independent given the state of information  $S$  if  $\{x,y|S\} = \{x|S\} \{y|S\}$  or equivalently, if  $\{x|y,S\} = \{x|S\}$ .

### b. Expansion

Regardless of whether  $x$  and  $y$  are probabilistically independent, we can write

$$\begin{aligned}\{x,y|S\} &= \{x|y,S\} \{y|S\} \\ &= \{y|x,S\} \{x|S\}\end{aligned}$$

We call this the "chain rule of probabilities". Note that for three events there are six possible representations.

$$\begin{aligned}\{x,y,z|S\} &= \{x|y,z,S\} \{y|z,S\} \{z|S\} \\ &= \{x|y,z,S\} \{z|y,S\} \{y|S\} \\ &= \{y|x,z,S\} \{x|z,S\} \{z|S\} \\ &= \{y|x,z,S\} \{z|x,S\} \{x|S\} \\ &= \{z|x,y,S\} \{x|y,S\} \{y|S\} \\ &= \{z|x,y,S\} \{y|x,S\} \{x|S\}\end{aligned}$$

For  $n$  variables there are  $n!$  possible expansions, each requiring the assignment of a different set of probabilities and each logically equivalent to the rest. However, while the assessments are logically equivalent there may be considerable differences in the ease with which the decision maker can provide them. Thus the question of which expansion to use in a problem is far from trivial.

#### c. Probability Trees

Associated with each expansion is a probability tree. Thus the expansion

$$\{x, y, z | S\} = \{x | y, z, S\} \{y | z, S\} \{z | S\}$$

implies the tree shown in Figure 5.16. The tree is a succession of nodes with branches emanating from each node to represent different possible values of a variable. The first assignment made is the probability of various values of  $z$ . The probability of each value of  $y$  is assigned conditioned on a particular value of  $z$ , and placed on the portion of the tree indicated by that value. Finally, the probabilities of various levels of  $x$  are assessed given particular values of  $z$  and  $y$  and placed on the portion of the tree specified by those values. When this has been done for all possible values of  $x$ ,  $y$ , and  $z$  the tree is complete. The probability of any particular path through the tree is obtained by multiplying the values along the branches and is  $\{x, y, z | S\}$ . Notice that the tree convention uses small circles to represent chance nodes. If we wish to focus on the succession in the tree rather than the detailed connections, we can draw the tree in the generic form shown in Figure 5.17.

#### d. Decision Trees

If a variable is controlled by a decision maker, it is represented in a tree by a decision node. Thus if  $y$  were a decision variable, Figure 5.17 could be redrawn as Figure 5.18. This tree states that the decision maker is initially uncertain about  $z$  and has assigned a probability distribution  $\{z | S\}$  to it. However, he will know  $z$  at the time he must set  $y$ , the decision variable. This node is represented, like all decision nodes, by a small square box. Once  $z$  and  $y$  are




FIGURE 5.16 A PROBABILITY TREE




FIGURE 5.17 A GENERIC PROBABILITY TREE




FIGURE 5.18 A GENERIC DECISION TREE

given, the decision maker will still be uncertain about  $x$ ; he has represented this uncertainty by  $\{x|y, z, S\}$ . Notice that a decision tree implies both a particular expansion of the probability assessments and a statement of the information available when a decision is made.

e. Probability Assignment for Decision Trees

The major problem with decision trees arises from the first of these characteristics. The order of expansion required by the decision tree is rarely the natural order in which to assess the decision maker's information. The decision tree order is only the simplest form for assessment when each variable is probabilistically dependent on all preceding aleatory and decision variables. If, as is usually the case, many independence assertions can be made, assessments are best done in a different order from that used in the decision tree. This usually means that we first draw a probability tree in an expansion form convenient to the decision maker and have him use this tree for assignment; it is called a probability assignment tree. Later the information is processed into the form required by the decision tree by representing it in one of the alternative expansion orders. This is often called "using Bayes' Rule" or "flipping the tree." It is a fundamental operation permitted by the arbitrariness in the expansion order.

Consider, for example, the decision tree of Figure 5.18 with one additional aleatory variable  $v$  added, as shown in Figure 5.19.

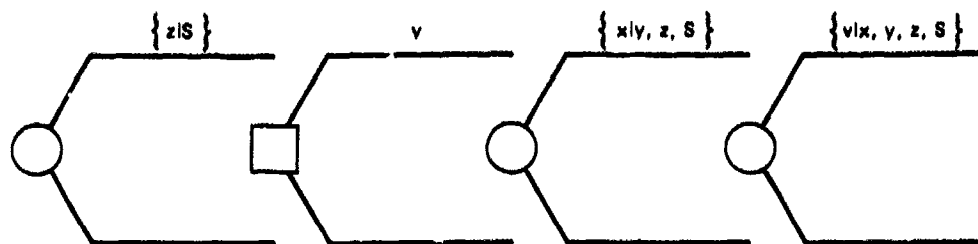


FIGURE 5.19 A FOUR NODE DECISION TREE

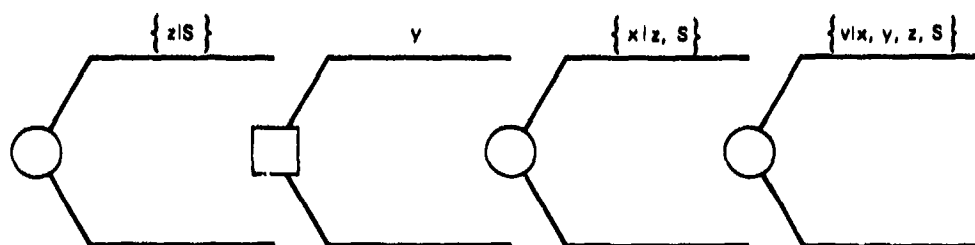


FIGURE 5.20 A FOUR NODE DECISION TREE GIVEN THE ASSERTION THAT  $y$  WILL NOT AFFECT  $x$




FIGURE 5.21 THE PROBABILITY ASSIGNMENT TREE

We interpret  $z$  as a test result that will become known,  $y$  as our decision,  $x$  as the outcome variable to which the test is relevant, and  $v$  as the value we shall receive if the test indicates  $z$ , we decide  $y$ , and  $x$  is the value of the outcome variable. Often  $y$  will not affect  $x$  in any way, even though  $y$  affects  $v$ . We write  $\{x|y,z,S\} = \{x|z,S\}$  to represent this assertion. Furthermore, it might be the case that the value will not depend on the test outcome but only on  $x$ . This assertion of independence is written:

$$\{v|x,y,z,S\} = \{v|x,y,S\}$$

With these independence assertions we have the tree shown in Figure 5.20. This tree requires the specification of  $\{z|S\}$  and  $\{x|z,S\}$ : the probability of various test results and the probability of various outcomes given test results. But typically in situations of this kind, the decision maker would prefer to assign directly the probabilities of different outcomes  $\{x|S\}$  and then the probabilities of differing test results given the outcome,  $\{z|x,S\}$ . In other words, he would prefer to make his assessments in the probability tree of Figure 5.21 and then have them processed to fit the decision tree of Figure 5.20. Since  $\{x|S\} \{z|x,S\} = \{z|S\} \{x|z,S\} = \{x,z|S\}$  this is no more than choosing one expansion over the other. The exact processing required for the decision tree is then the summation,

$$\{z|S\} = \sum_x \{z|x,S\} \{x|S\}$$

and division:

$$\{x|z,S\} = \frac{\{z|x,S\} \{x|S\}}{\{z|S\}}$$

Recall, however, that this whole procedure was possible only because variable  $x$  did not depend on the decision variable  $y$ .

#### f. Influence Diagrams

An influence diagram is a way of describing the dependencies among random variables and decisions. An influence diagram can be used to visualize the probabilistic dependencies in a decision analysis,



and specify the states of information for which independencies can be assumed to exist.

Figure 5.22 shows how influence diagrams represent the dependencies among random variables and decisions. A random variable

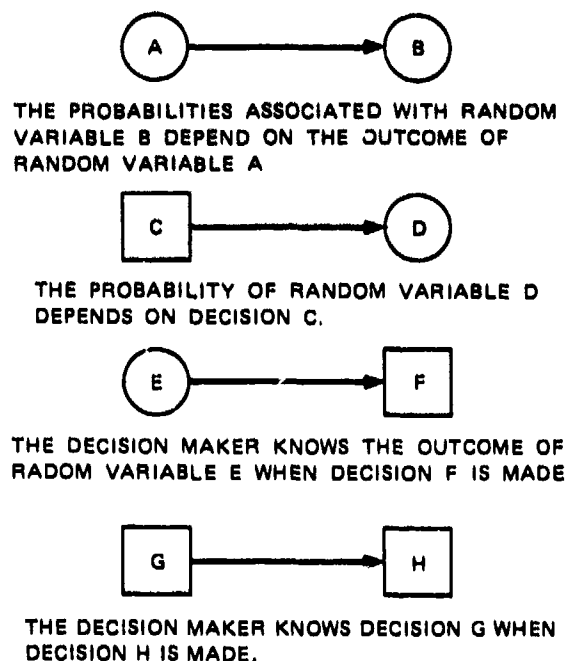
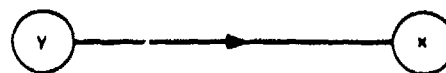


FIGURE 5.22 DEFINITIONS USED IN INFLUENCE DIAGRAMS

is represented by a circle containing its name or number. An arrow pointing from random variable A to random variable B means that the outcome of A can influence the probabilities associated with B. An arrow pointing to a decision from either another decision or a random variable means that the decision is made with the knowledge of the outcome of the other decision or random variable. A connected set of squares and circles is called an influence diagram because it shows how random variables and decisions influence each other.

The influence diagram in Figure 5.23(a) states that the probability distribution assigned to  $x$  may depend on the value of  $y$ , while the influence diagram in Figure 5.23(b) asserts that  $x$  and  $y$  are probabilistically independent for the state of information with which



(a) A SIMPLE INFLUENCE DIAGRAM



(b) AN EVEN SIMPLER INFLUENCE DIAGRAM



(c) AN ALTERNATE REPRESENTATION

FIGURE 5.23 TWO NODE INFLUENCE DIAGRAMS

the diagram was drawn. Note that the diagram of Figure 5.23(a) really makes no assertion about the probabilistic relationship of  $x$  and  $y$  since, as we know, any joint probability  $\{x, y|S\}$  can be represented in the form  $\{x, y|S\} = \{x|y, S\} \{y|S\}$ . However, since  $\{x, y|S\} = \{y|x, S\} \{x|S\}$ , the influence diagram of Figure 5.23(a) can be redrawn as shown in Figure 5.23(c); both are completely general representations requiring no independence assertions. While the direction of the arrow is irrelevant for this simple example, it is used in more complicated problems to specify the states of information upon which independence assertions are made.

Similarly, with three variables  $x$ ,  $y$ ,  $z$  there are six possible influence diagrams of complete generality, one corresponding to each of the possible expansions we developed earlier. They are shown in Figure 5.24. While all of these representations are logically equivalent, they again differ in their suitability for assessment purposes. In large decision problems, the influence diagrams can display the needed assessments in a very useful way.




FIGURE 5.24 ALTERNATE INFLUENCE DIAGRAMS FOR  $\{x, y, z\}$

g. Graphical Manipulation

Since there are many alternative representations of an influence diagram, we might ask what manipulations can be performed on an influence diagram to change it into another form that is logically equivalent.


The first observation we should make is that an arrow can always be added between two nodes without making an additional assertion

about the independence of the two corresponding variables. That is, saying that  $x$  may depend on variable  $y$  is not equivalent to saying that  $x$  must depend on  $y$ . Thus the diagram of Figure 5.23(b) can be changed into either of the diagrams shown in Figures 5.23(a) and 5.23(c) without making an erroneous assertion. However, the reverse procedure could lead to an erroneous assertion. Creating additional influence arrows will not change any probability assessment, but may destroy explicit recognition of independencies in the influence diagram.


Thus, Figures 5.23(a) and 5.23(c) are two equivalent influence diagrams. They are equivalent in that they imply the same possibility of dependencies between  $x$  and  $y$  given the state of information on which the diagram was based.

An arrow joining two nodes in an influence diagram may be reversed provided that all probability assignments are based on the same set of information. For example, consider the influence diagram of Figure 5.25(a). Since the probability assignments to both  $x$  and  $y$  are made given knowledge of  $z$  the arrow joining them can be reversed as shown in Figure 5.25(b) without making any incorrect or additional assertions about the possible independence of  $x$  and  $y$ . Figure 5.25(c) shows another example where the assignment of probability to  $x$  does not depend on the value of  $z$ , and so it might appear that no reversal was possible. However, recall that we can always add an arrow to a diagram without making an incorrect assertion. Thus we can change the diagram of Figure 5.25(c) to that of Figure 5.25(a), and then that of Figure 5.25(a) to that of Figure 5.25(b). The influence arrow between  $y$  and  $x$  can be reversed after an influence arrow is inserted between  $z$  and  $x$ .


The graphical manipulation procedure may yield more than one result. For example, consider the reversal of the three-node influence diagram shown in Step 1 of Figure 5.26(a). Suppose we first attempt to reverse the  $y$  to  $x$  arrow. In order that  $x$  and  $y$  have only common influences, we must provide  $x$  with an influence from  $z$  (Step 2), before performing the reversal (Step 3). Since both  $x$  and  $y$  now are conditioned on only  $z$ , the influence joining them may be reversed



(a) AN INFLUENCE DIAGRAM



(b) ARROW BETWEEN x AND y REVERSED



(c) ANOTHER INFLUENCE DIAGRAM

FIGURE 5.25 GRAPHICAL MANIPULATION OF INFLUENCE DIAGRAMS

(Step 4). Finally, since both  $x$  and  $y$  are assigned probabilities after  $z$  is known, the influence joining them can be reversed (Step 5).

Suppose, however, that the same diagram (Step A, Figure 5.26(b)) were transformed by first reversing the arrow joining  $z$  and  $y$  (Step B), which is possible since  $y$  and  $z$  are based on the same state of information (i.e. there are no arrows into  $y$  or  $z$  from any other node in Step A of Figure 5.26(b)). Then the arrow joining  $x$  and  $y$  could be reversed (Step C) because neither  $x$  nor  $y$  now has an arrow leading into it from any other node. Both this transformation and the one in Figure 5.26(a) are correct. However, Step C of Figure

5.26(b) shows that there is no need to indicate conditioning of  $z$  on  $x$ . Step 5 of Figure 5.26(a) contains this unnecessary but not incorrect influence.

#### h. Influence Diagrams with Decision Variables

We shall now extend the concept of influence diagrams to include decision variables. We begin with a formal definition of influence diagrams.

An influence diagram is a directed graph having no loops. It contains two types of nodes:

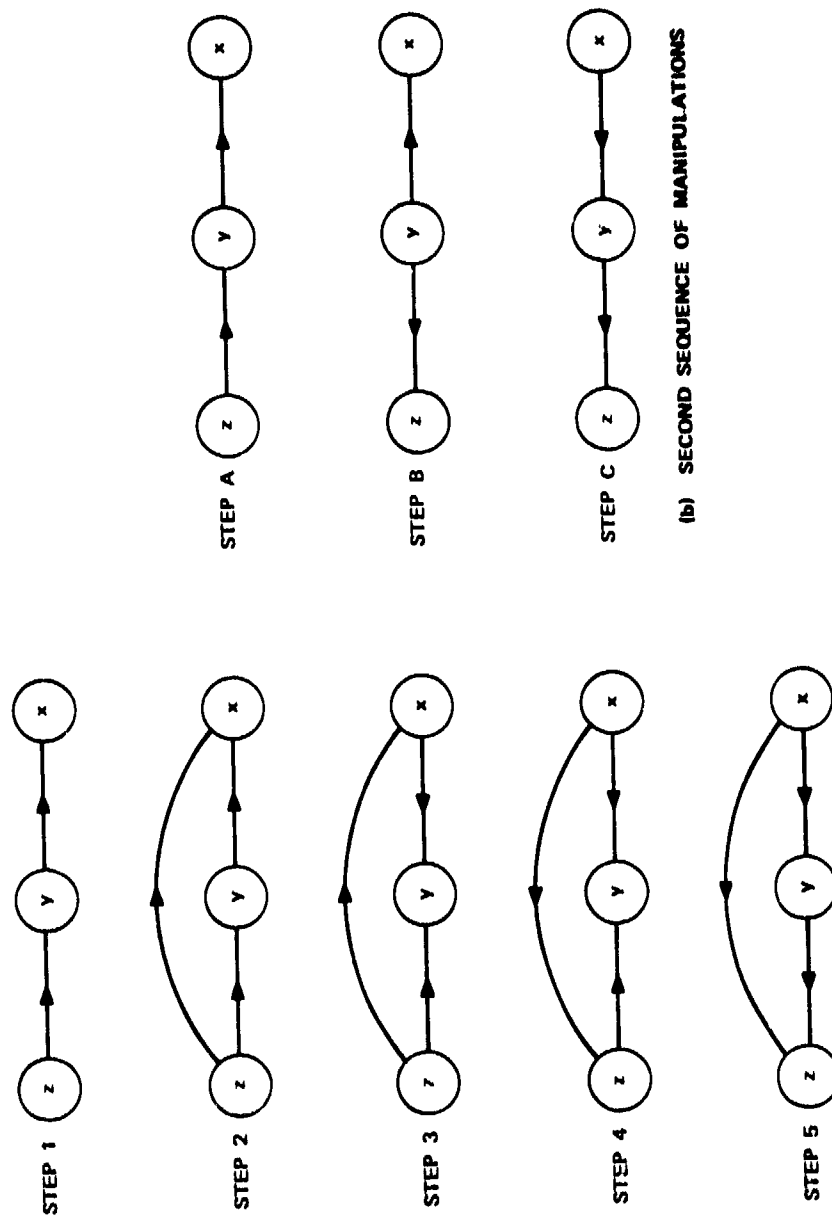
- Decision nodes represented by boxes ( );
- Chance nodes represented by circles ( ).

Arrows between node pairs indicate influences of two types:

- Informational influences represented by arrows leading into a decision node, these show exactly which variables will be known by the decision maker at the time that the decision is made;
- Conditioning influences, represented by arrows leading into a chance node. These show the variables on which the probability assignment to the chance node variable will be conditioned.

The informational influences on a decision node represent a basic cause-effect ordering whereas the conditional influences into a chance node represent, as we have seen, a somewhat arbitrary order of conditioning which may not correspond to any cause-effect notion and which may be changed by application of the laws of probability (e.g. Bayes' Rule).

Figure 5.27 is an example of an influence diagram. Chance node variables  $a, b, c, e, f, g, h, i, j, k, l, m,$  and  $o$  all indicate aleatory variables whose probabilities must be assigned given their respective conditioning influences. Decision node variables  $d$  and  $n$  represent decision variables that must be set as a function of their respective informational influences. For example, the probability assignment to variable  $i$  is conditioned upon variables  $f, g,$  and  $l$ , and only these variables. In inferential notation



(a) FIRST SEQUENCE OF MANIPULATIONS

FIGURE 5.26 GRAPHICAL MANIPULATIONS PRODUCING NON-UNIQUE RESULTS




FIGURE 5.27 AN INFLUENCE DIAGRAM WITH DECISION NODES

this assignment is  $\{i|f,g,l,E\}$ , where  $E$  represents a special  $S$ , the initial state of information upon which the construction of the entire diagram is based. As another example, the decision variable  $d$  is set with knowledge of variables  $a$  and  $c$ , and only these variables. Thus,  $d$  is a function of  $a$  and  $c$ .

#### 1. Node Terminology

One of the most important, but most subtle, aspects of an influence diagram is the set of possible additional influences that are not shown on the diagram. An influence diagram asserts that these missing influences do not exist.

In order to illustrate this characteristic of influence diagrams more clearly we must make a few more definitions.

A path from one node to another node is a set of influence arrows connected head to tail that forms a directed line from one node to another.

With respect to any given node we make the following definitions:

The predecessor set of a node is the set of all nodes having a path leading to the given node.



The direct predecessor set of a node is the set of nodes having an influence arrow connected directly to the given node.

The indirect predecessor set of a node is the set formed by removing from its predecessor set all elements of its direct predecessor set.

The successor set of a node is the set of all nodes having a path leading from the given node.

The direct successor set of a node is the set of nodes having an influence arrow connected directly from the given node.

The indirect successor set of a node is the set formed by removing from its successor set all elements of its direct successor set.

We refer to members of these sets as predecessors, direct predecessors, indirect predecessors, successors, direct successors, and indirect successors. Figure 5.28 shows the composition of each of these sets in relation to node g.

#### j. Missing Influences

We now are prepared to investigate the implications of influences not shown in a diagram. A given node could not have any arrows coming into it from successor nodes because this addition would form a prohibited loop in the diagram. However, it could conceivably have an additional arrow coming from any other node.

The situation for decision nodes is relatively simple. The diagram asserts that the only information available when any decision is made is that represented by the direct predecessors of the decision. The addition of a new arrow, or informational influence, would usually add to the information available for decision making, and destroy the original logic of the diagram. The influence diagram asserts that this information

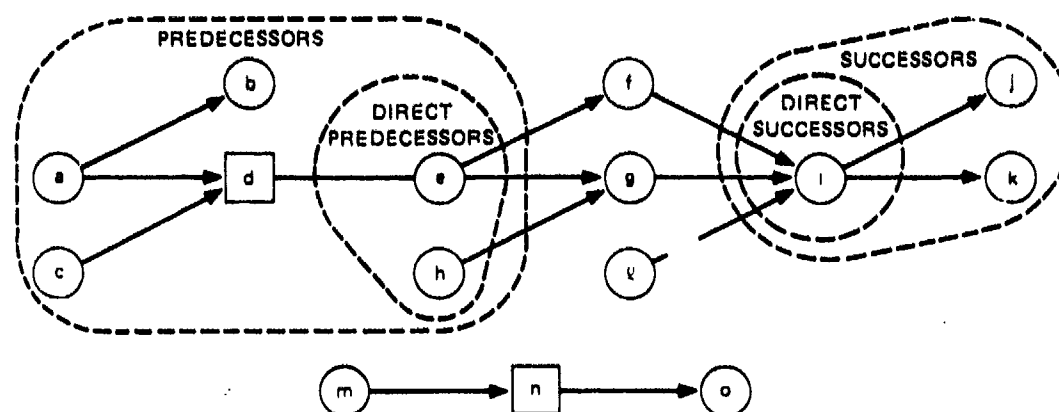


FIGURE 5.28 SOME SETS DEFINED BY THE NODE  $g$

is not directly available; however, all or part of it might be inferred indirectly from the direct predecessor set.

The situation for chance nodes is more complex. The diagram partially constrains the probability conditioning (expansion) order for chance nodes. In general, the probability assignment for a given chance node,  $x$ , might be conditioned on all non-successors (except for  $x$  itself). Let us call this set  $N_x$ , and let  $D_x$  be the set of direct predecessors of  $x$ . The set  $D_x$  is, of course, contained in  $N_x$ . The diagram asserts that the probability assignment to  $x$  given  $N_x$  is the same as to  $x$  given  $D_x$ ; that is,

$$\langle x | N_x, E \rangle = \langle x | D_x, E \rangle$$

The addition of a new arrow or conditioning influence from an element of  $N_x$  to  $x$  would increase the set of direct predecessors and seem to increase the dimensionality of the conditional probability assignment. While this addition would not violate the logic of the diagram, it would cause a loss of information regarding independence of the added conditioning influence. The original diagram asserts that all information in the set  $N_x$  that is relevant to the probability assignment to  $x$  is indirectly summarized by the direct predecessors  $D_x$ . In classical terms with respect to  $x$ ,  $D_x$  is a sufficient statistic for  $N_x$ .

Returning to Figure 5.27 as an example, the probability assignment to variable  $g$  is in principle conditioned on all variables except  $g$ ,  $i$ ,  $j$ , and  $k$ . However, the diagram asserts that the variables on which  $g$  depends are sufficiently summarized by only  $e$  and  $h$ . This means  $\{g|a,b,c,d,e,f,h,i,m,n,o,E\} = \{g|e,h,E\}$ . This is a strong and useful assertion relating many of the variables in the diagram by the lack of arrows as well as by the ones that are present.

We have seen that an influence diagram indicates a specific, but possibly non-unique order for conditioning probability assignments as well as the information available as the basis for each decision. When decision rules are specified for each decision node and probability assignments are made for each chance node, the influence diagram relationships can be used to develop the joint probability distribution for all variables.

#### k. Relationship of Influence Diagrams to Decision Trees

Some influence diagrams do not have corresponding decision trees. As in a decision tree, all probability assignments in an influence diagram--including the assignment limitations represented by the structure--must be based on a base state of information,  $E$ . Unlike a decision tree, the nodes in an influence diagram do not have to be totally ordered nor do they have to depend directly on all predecessors. The freedom from total ordering allows convenient probabilistic assessment and computation. The freedom from dependence on all predecessors allows decisions to be based on informational event sets that are incompatible with a "single decision maker" point of view. If a single decision maker is assumed not to forget information, then the direct predecessor set of one decision must be a subset of the direct predecessor set of any subsequent decision. In the influence diagram of Figure 5.28 decisions  $d$  and  $n$  have mutually exclusive direct predecessor sets,  $(a,c)$  and  $(m)$ . This situation could not be represented with a decision tree.

If the informational arrows shown as dashed lines in Figure 5.29 are added to Figure 5.28, then the influence diagram can be represented by a decision tree. Many different valid decision trees can be

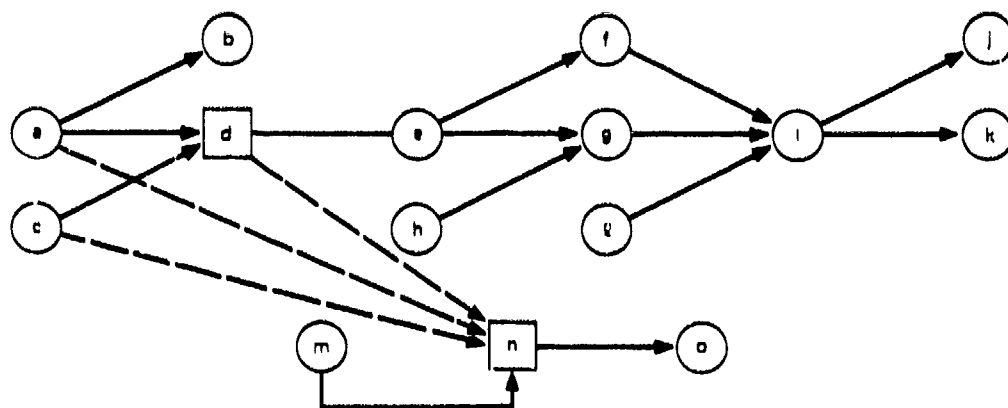


FIGURE 5.29 AN INFLUENCE DIAGRAM REPRESENTABLE BY A DECISION TREE

constructed from this new influence diagram. The only conditions are that they must (1) preserve the ordering of the influence diagram and (2) not allow a chance node to be a predecessor of a decision node for which it is not a direct predecessor. For example, the chance node  $m$  must not appear ahead of decision node  $d$  in a decision tree because this would imply that the decision rule for  $d$  could depend on  $m$ , which is not the case.

The situation becomes more complex when we add a node such as  $p$  in Figure 5.30. If we were to construct a decision tree beginning

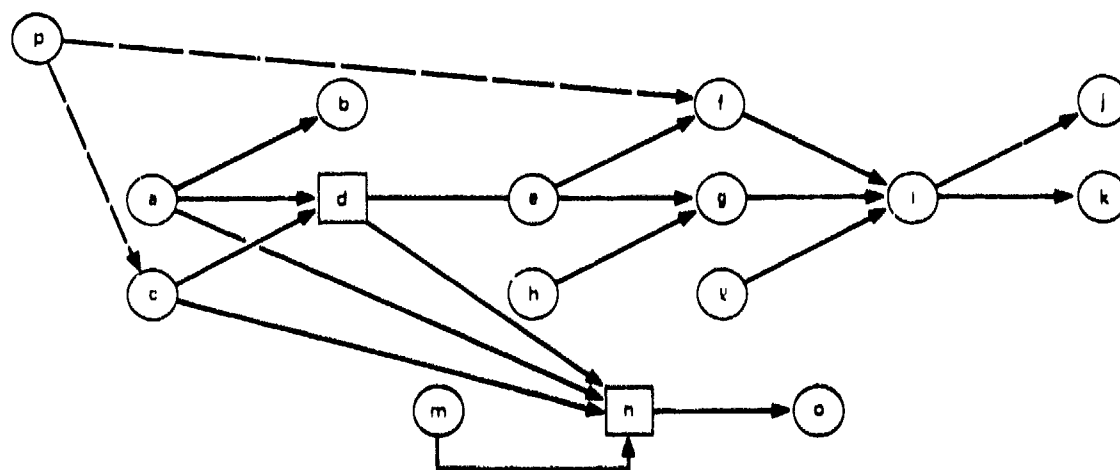


FIGURE 5.30 INFLUENCE DIAGRAM REQUIRING PROBABILISTIC MANIPULATION BEFORE DECISION TREE CONSTRUCTION

with chance node  $p$  it would imply that the decision rules at nodes  $d$  and  $n$  could depend on  $p$ , which is not the actual case. Node  $p$  represents a variable that is used in the probability assignment model but that is not observable by the decision maker at the time that he makes his decisions. In this situation, we would normally use the laws of probability (e.g. Bayes' Rule) to eliminate the conditioning of  $c$  on  $p$ . This process would lead to a new influence diagram reflecting a change in the sequence of conditioning. This would result in the inclusion of additional influences.

In Figure 5.31, the dashed arrow represents an influence that has been "turned around" by Bayes' Rule. The resulting diagram can




FIGURE 5.31 INFLUENCE DIAGRAM READY FOR DEVELOPMENT INTO A DECISION TREE

be developed into a decision tree without further processing of probabilities. Also note that the change in the influence diagram required only information already specified by the original influence diagram (Figure 5.30) and its associated numerical probability assignments. Thus it can be carried out by a routine procedure.

The foregoing considerations motivate two new definitions.

A decision network is an influence diagram:

- (i) that implies a total ordering among decision nodes,
- (ii) where each decision node and its direct predecessors directly influence all successor decision nodes.

A decision tree network is a decision network:

- (iii) where all predecessors of each decision node are direct predecessors.

Requirement (i) is the "single decision maker" condition and requirement (ii) is the "no forgetting" condition. These two conditions guarantee that a decision tree can be constructed, possibly after some probabilistic processing. Requirement (iii) assures that no probabilistic processing is needed so that a decision tree can be constructed in direct correspondence with the influence diagram.

As an example consider the standard inferential decision problem represented by the decision network of Figure 5.32(a). As discussed earlier, this influence diagram cannot be used to generate a decision tree directly because the decision node *c* has a non-direct predecessor that represents an unobservable chance variable. To convert this decision network to a suitable decision tree network we simply reverse the arrow from *a* to *b*, which is permissible because they have only common predecessors, namely none. We thus achieve the decision tree network of Figure 5.32(b), and with redrawing we arrive at Figure 5.32(c).

Specifying the limitations on possible conditioning by drawing the influence diagram may be the most significant step in probability assignment. The remaining task is to specify the numerical probability of each chance node variable conditioned on its direct predecessor variable. This task can be carried out using the assessment procedure discussed in Chapter IV.

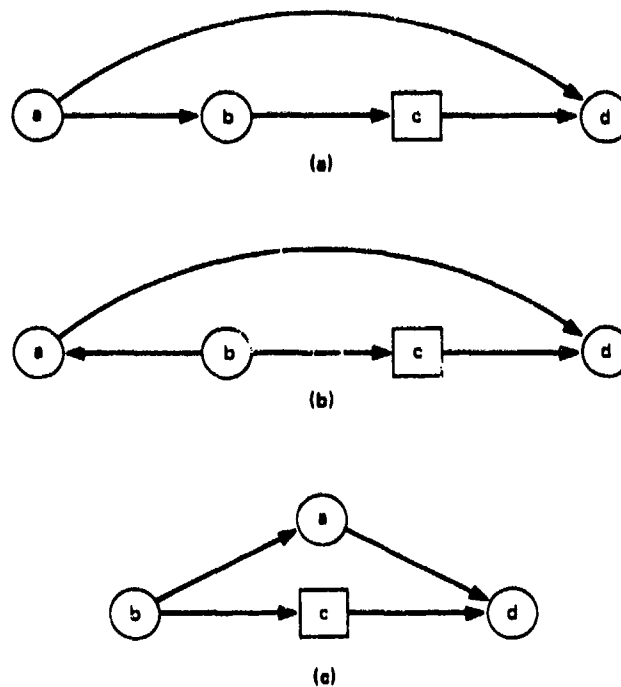


FIGURE 5.32 THE PROCESS OF CONVERTING A DECISION NETWORK TO A DECISION TREE NETWORK

## 2. An Example of the Use of Influence Diagrams

The importance of improved methods of problem modeling can best be appreciated in an example. The example that follows is a modified version of an analysis performed a few years ago on the value of information gathering with respect to U.S. policy decisions regarding the Persian Gulf.

Certain recent events have had a major effect on the timeliness of the analysis. We shall present the example from the viewpoint of the time in it was prepared rather than rewrite it to incorporate what time has revealed.

a. Description of a Model Based on Influence Diagrams

The approach used in the analysis was to model the sequence of events that might lead from U.S. policy decisions to the development of potential conflicts in the Persian Gulf. U.S. policy decisions do not always have a direct effect on the likelihood that a conflict will occur in the Persian Gulf. More often U.S. actions will influence one or more of a set of interrelated and uncertain events, which in turn affect the likelihood of conflict.

The model required subjective assessments of the probabilities that uncertain events would occur in the Persian Gulf in the subsequent five year period. However the probability that any one event might occur was influenced by whether or not other events had already occurred. Thus, the person making the assessment had to condition his probability estimates on assumptions about the occurrence or non-occurrence of other events. For example, an estimate of the likelihood that a revolutionary regime would take power in Saudi Arabia within the five-year period depended on whether or not King Faisal died and whether or not the traditional government of Saudi Arabia avoided political infighting and instability.

The model did not require a subject to estimate the time at which an uncertain event would occur. Instead he had to estimate the probability that the event would happen within the following five years, or before the occurrence of a major war in the area.

Because the model was concerned primarily with events that could be influenced by U.S. policy decisions, not all the events that could affect the likelihood of a conflict in the Persian Gulf were included in the model. In fact, in some cases, events not included in the model might have influenced events in the Gulf more strongly than those shown in the model. For example, the occurrence of a revolution in Saudi Arabia during the five year period may have been influenced more by the rate of Saudi Arabian economic and social development than by the level of revolutionary activity in the Persian Gulf. However, there appeared to be few, if any, U.S. policy decisions that could directly influence the rate of political and social development in Saudi Arabia. The effects of events



not shown explicitly in the model were reflected, however, in the probability assessments.

Table 5.1 shows the sequence of decisions and events that were actually used. Since U.S. decisions affecting the Persian Gulf might have been based on information derived from several different intelligence activities, the first step in the model was to hypothesize the deployment of a particular intelligence activity. For the purpose of the analysis, it was assumed that an intelligence activity would gather information about one or more of the events described in the model. If an intelligence activity was deployed, the United States would receive information from this activity, predicting whether that event would or would not occur.

After receiving this information from the intelligence activity, the United States would make some initial policy decisions with respect to the Persian Gulf. The alternatives available, which would influence the general course of events in the Persian Gulf, included continuing strong support for Israel and removing the U.S. naval force stationed in the Persian Gulf.

While the United States made these decisions, the Soviet Union would be making a similar set of decisions that would influence the course of events in the Persian Gulf. From the point of view of U.S. policy, Soviet decisions appeared as uncertain events. Soviet influence was characterized in the model by two principal decisions: whether or not to increase the flow of military equipment to revolutionary organizations in the area, and whether or not to introduce a naval force comparable to the U.S. force in the Persian Gulf.

Clearly, these two choices were an abstraction of the many possibilities available to the Soviet Union. For example, the Soviet Union could indirectly but substantially increase the amount of military equipment supplied to revolutionary organizations by increasing arms shipments to either Iraq or South Yemen. Similarly, in choosing to introduce naval forces into the Persian Gulf, the Russians had several options as

Table 5.1

SEQUENCE OF DECISIONS AND EVENTS  
IN THE INFLUENCE DIAGRAM EXAMPLE

- I. U.S. decision to deploy and intelligence activity
- II. Information received from the intelligence activity
  1. The activity reports that an event will occur.
- III. Initial U.S. policy decisions
  2. The United States continues its strong diplomatic support for Israel and continues to supply significant amounts of military equipment to Israel.
  3. The United States moves the naval force it currently has in the Persian Gulf to a base in the Indian Ocean.
- IV. Soviet Influence
  4. The Soviets substantially increase the amount of military equipment they supply to revolutionary organizations in the Persian Gulf.
  5. The Soviets increase their own military presence in the Persian Gulf.
- V. Arab-Israeli conflict and revolutionary activity
  6. The Arabs and Israel reach a political settlement that is acceptable to most Arabs in the Persian Gulf states.
  7. There is a significant increase in the level of revolutionary organization and activity in the Persian Gulf states.
- VI. Instability and revolution in Saudia Arabia
  8. King Faisal dies.
  9. Considerable internal instability develops in Saudi Arabia
- VII. Instability and revolution in Iran
  11. The Shah of Iran dies.
  12. Considerable internal instability develops in Iran.
  13. A revolutionary regime takes power in Iran.

Table 5.1 (Concluded)

- VIII. U.S. policy decisions to support Persian Gulf states
  - 14. The United States gives diplomatic support and large amounts of military equipment to Saudi Arabia.
  - 15. The United States gives diplomatic support and large amounts of military equipment to Iran.
- IX. Conflict between Saudi Arabia and Iran
  - 16. A revolutionary regime takes power in Qatar, Bahrain, or the U.A.E.
  - 17. Saudi Arabia and Iran engage in a serious military conflict.
  - 18. Iraq enters the conflict between Saudi Arabia and Iran.
- X. Conflict between Iraq and Kuwait or Iran.
  - 19. A revolutionary regime takes power in Kuwait.
  - 20. Iraq engages in military conflict with Kuwait--possibly in an attempt to annex Kuwait.
  - 21. Iraq and Iran enter a serious military conflict.
  - 22. Saudi Arabia enters the conflict between Iran and Iraq.

to the composition and strength of those forces. However, the Soviet decision hypothesized in the model served as a first approximation of possible Soviet influence in the area.

Once the decisions of the United States and Soviet Union had been specified, it was possible to estimate the likelihood of events that set the stage for potential conflicts in the Persian Gulf. The two events considered in this phase of the model were the future development of the Arab-Israeli conflict and the level of revolutionary organization and activity in the Persian Gulf states. These two events were closely interrelated, and this relationship was reflected in the model.

After the status of the Arab-Israeli conflict and the level of revolutionary activity were specified, the model considered the possibility of instability and revolution in both Saudi Arabia and Iran, the two major traditional powers in the area. For each country, the model began with the possibility that the current national leader would die within the next five years--an important factor given the individualistic style of government in both countries. The next step was to consider the likelihood that the governments of Saudi Arabia and Iran would experience a period of internal instability and political turmoil, conditioned on whether or not the current rulers of each country were living. Then, depending on whether or not these events had occurred, the model dealt with the likelihood of revolutionary changes in government in either Saudi Arabia or Iran or both.

Once the types of government in Saudi Arabia and Iran had been determined, the United States had to decide whether or not to continue support for those Persian Gulf states. To determine whether the type of governments in Saudi Arabia and Iran had an influence on U.S. policy, the model considered the U.S. support decisions after resolving the question of revolutions. In particular, the model addressed the question of whether it was a good idea for the United States to continue its support for Saudi Arabia or Iran after a revolutionary government had taken power.

When the U.S. support decisions had been made, the model specified all the events and decisions needed to allow estimates of the

likelihood of major conflicts in the Persian Gulf. In a sense, the steps numbered I to VIII in Table 5.1 defined the various scenarios, or alternative sequences of events and decisions, that might lead to a major conflict in the Persian Gulf. For each of these possible scenarios, the model considered the possibility of a major conflict between Saudi Arabia and Iran, and the most likely ways in which such a conflict could develop. The model also allowed for the possibility that Iraq might choose to enter into such a conflict.

Similarly, the model considered the possibility of a major conflict between either Iraq or Iran, or between Iraq and Kuwait. These two conflicts were closely related, since the Shah of Iran had stated that he would come to the aid of Kuwait in the event of an attempted Iraqi takeover of Kuwait. As before, the model allowed for the possibility that the Saudi Arabians might become involved in such a conflict between Iraq and Iran.

By combining the occurrence or nonoccurrence of each of the events and decisions in the model, it was possible to develop a large number of Persian Gulf scenarios, as demonstrated in Figure 5.33. The figure shows the scenarios in the form of a decision tree, starting on the left and developing to the right. Each node in the tree represents either a U.S. decision or an uncertain event. U.S. decisions are represented by squares and uncertain events are represented by circles. For example, the small square to the left of the "DEPLOY ACTIVITY" branch indicated that the deployment of an intelligence activity was a U.S. decision.

By progressing from the left to the right along any particular path through the tree, it was possible to specify a unique Persian Gulf scenario. For example, taking the branch marked "DEPLOY ACTIVITY" would mean that we were considering only those scenarios in which the United States had deployed an intelligence activity to gain information about one of the subsequent events in the model. If an intelligence activity was deployed, it could either predict that an event would occur or predict that an event would not occur. If we assumed that the activity had predicted that an event would occur, we would proceed to the branch

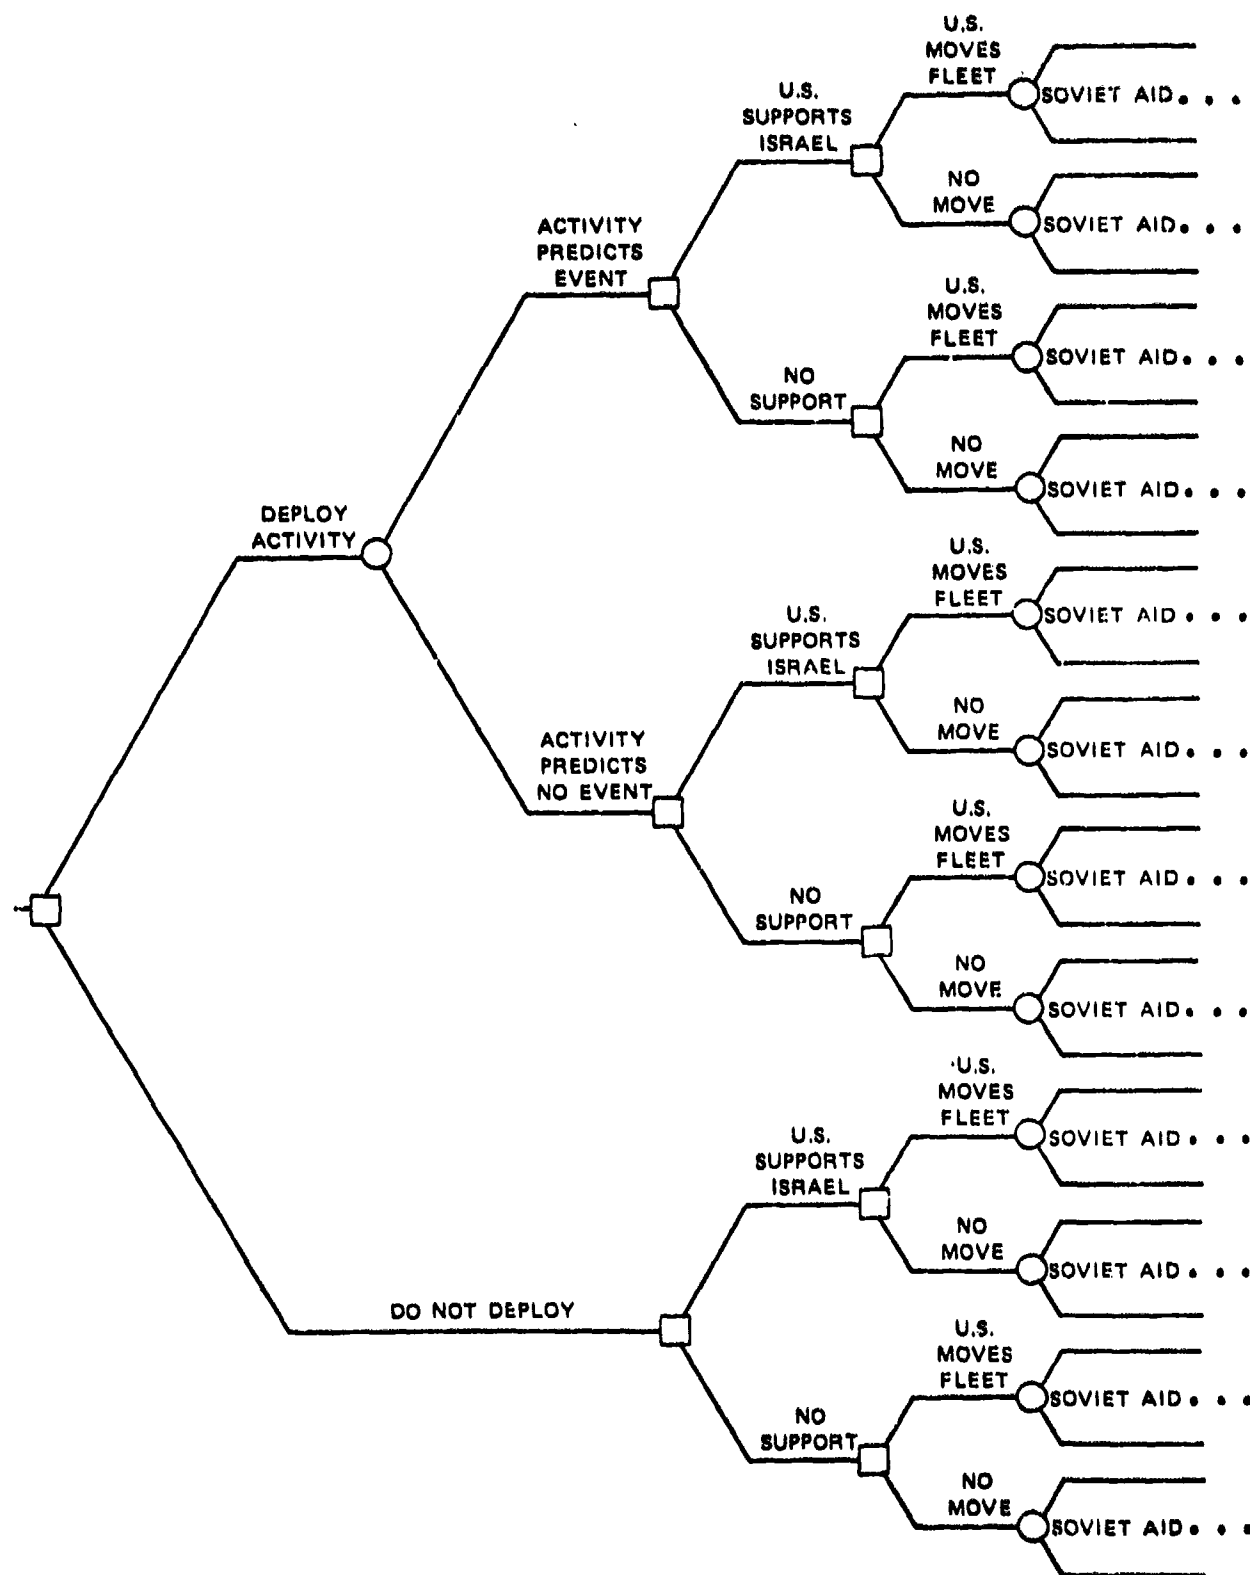


FIGURE 5.33 DEVELOPMENT OF PERSIAN GULF SCENARIOS

marked "ACTIVITY PREDICTS EVENT," which leads to another U.S. decision. The decision at that point was whether or not the United States should continue its strong support for Israel. Assuming that the United States decided to continue supporting Israel was equivalent to taking the path marked "U.S. SUPPORTS ISRAEL" near the top of Figure 5.33. At the end of this branch was another small square indicating another U.S. decision: whether or not to remove the naval force stationed in the Persian Gulf. A U.S. decision to move the fleet out of the Gulf was equivalent to moving down the branch marked "U.S. MOVES FLEET." At the end of this branch is an event node indicated by a small circle, which dealt with the level of Soviet aid to revolutionary activities and organizations in the Persian Gulf. Thus, by continuing to take one branch after another through the tree, we could completely specify one scenario for the Persian Gulf.

Other scenarios were specified by other paths through the tree. For example, if we took the tree branches closest to the bottom of Figure 5.33, we had the scenario where the United States did not deploy an intelligence activity, did not support Israel, and did not move its fleet out of the Persian Gulf.

When we added all of the branches associated with the many decisions and events in the model, the decision tree became so large that it was not possible to show the entire tree in a single figure. Figure 5.34 shows the generic decision tree representing the Persian Gulf scenarios. Here one must imagine that each of the disconnected decision or event nodes is attached to the ends of all of the preceding branches in the decision tree. Connecting the various decisions and event nodes in this way would generate the entire decision tree and specify all of the scenarios in the Persian Gulf model.

A simple calculation showed that there were approximately 6.3 million scenarios considered in this model. While many of these scenarios were very unlikely, each contributed to the conclusions of the model. Instead of considering only those scenarios that had the highest probability or seemed the most plausible--as one might in an intuitive analysis of the problem--the model allowed us to aggregate the effects of many low-probability scenarios.

Even though there were relatively few U.S. decisions represented in the Persian Gulf model, there were a large number of policies available to the United States. A policy was a complete description of U.S. decisions that would be taken in response to any possible future scenario. There were approximately 262,000 U.S. policies represented by the Persian Gulf model.

b. Using the Model

A model as large as the one for the Persian Gulf had two related problems. First, it was difficult to manipulate a model representing over 6 million scenarios, even using the capabilities of large, high-speed computers. Even when the computer was capable of handling this many scenarios, it was difficult to interpret the results without a more efficient way of conceptualizing the problem. The second problem was potentially even more limiting. If each of the millions of uncertain events represented in the model required a separate probability assessment, it would never have been feasible to collect the data necessary to analyze the model and study its implications.

The answer to these problems lay in carefully defining the dependencies among the uncertain events in the model and in using these dependencies to eliminate redundant portions of the model. While the decision tree shown in Figures 5.33 and 5.34 shows explicitly all of the scenarios in the model, it does not show which events and decisions had to be known in order to assess the likelihood of any uncertain event.

For example, when assessing the likelihood that the Arab-Israeli conflict would be resolved in the next five years, a middle East expert may have wished to know whether or not the United States would continue its strong military support for Israel. However, when assessing the probability that there would be an increase in the level of revolutionary activity in the Persian Gulf, the same expert might have been interested only in whether or not there was a settlement in the Arab-Israeli dispute, not in whether the settlement was brought about by a continuation of U.S. support for Israel. In this case, the likelihood of an increase






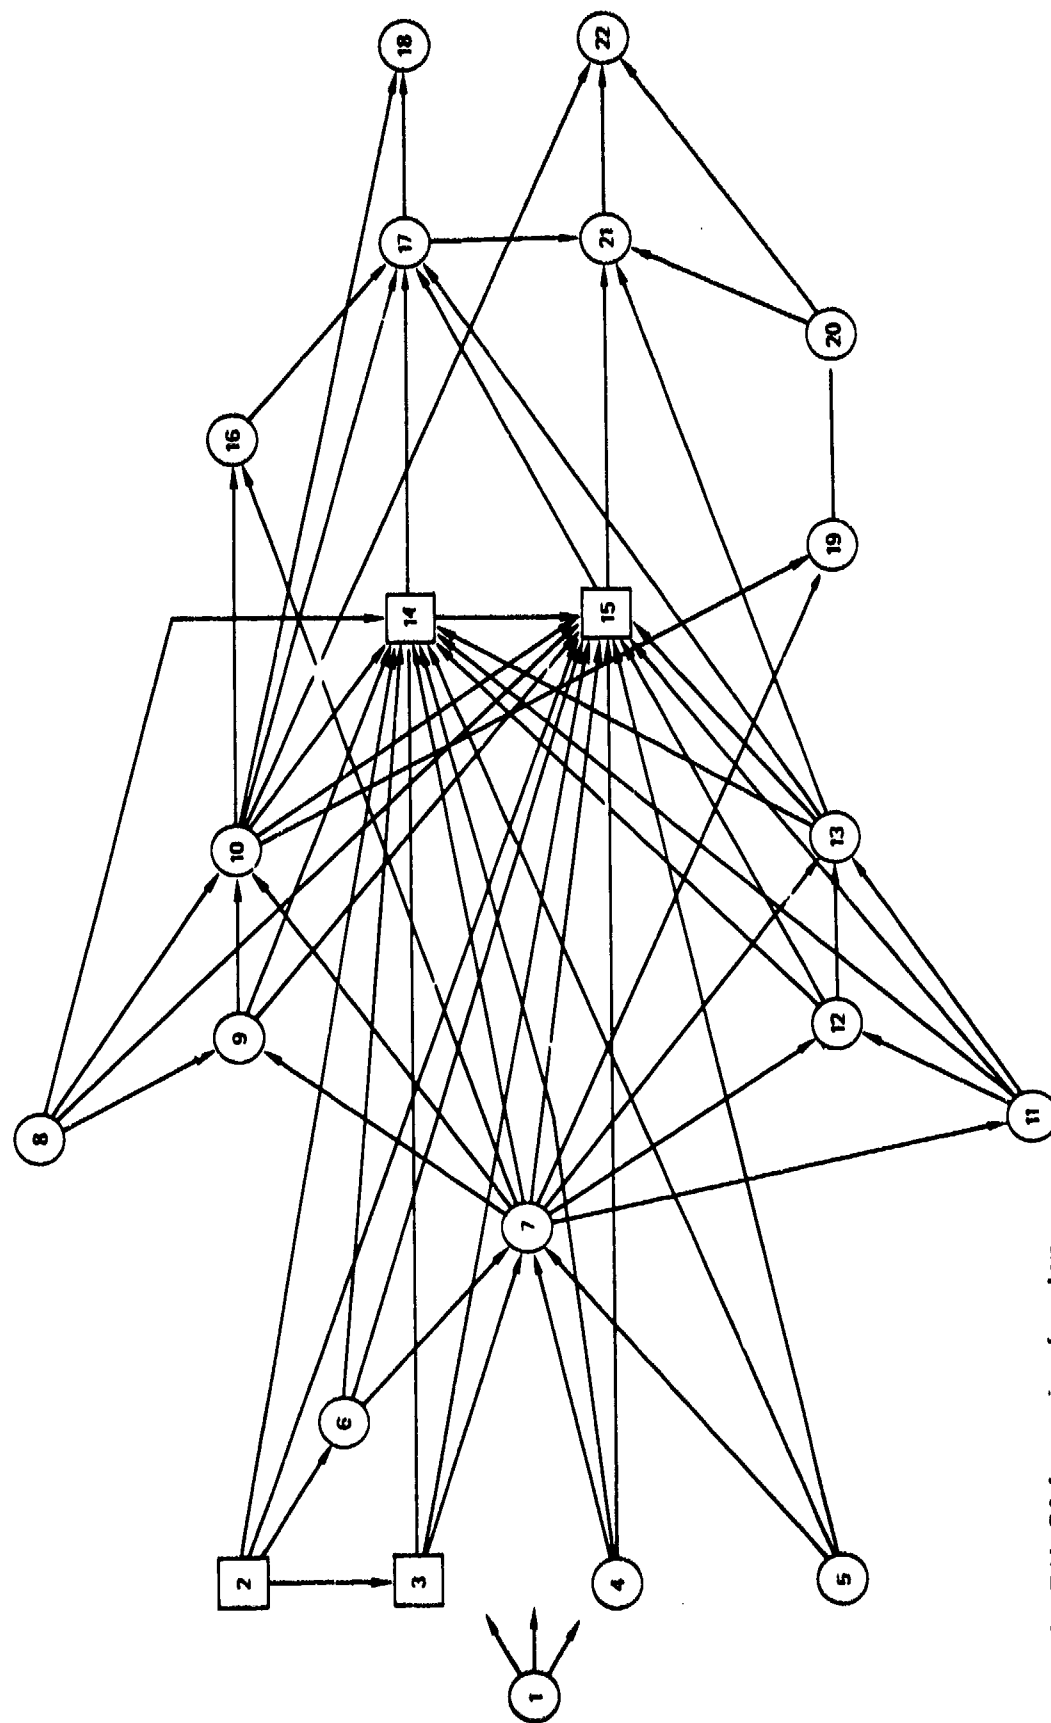
FIGURE 5.34 SIMPLIFIED DIAGRAM OF PERSIAN GULF SCENARIOS

in revolutionary activity was independent of U.S. policy once the status of the Arab-Israeli conflict had been specified.

To show such dependencies among uncertain events, we used an influence diagram. The influence diagram in Figure 5.35 shows the assumed dependencies of events and decisions in the Persian Gulf model. The numbers in the circles and squares corresponded to the numbered events and decisions in Table 5.1. For example, the square containing the number 2 represented the U.S. decision to continue giving Israel strong diplomatic support and significant amounts of military equipment. The circle containing a 6 represented the uncertain event that the Arabs and Israel would reach a political settlement that was acceptable to most of the Arabs in the Persian Gulf states. The arrow leading from decision 2 to event 6 meant that the probability that the Arabs and Israel would reach an acceptable political settlement depended on whether or not the United States continued its strong diplomatic support of Israel.

To estimate the likelihood of any uncertain event shown in Figure 5.35, it was necessary to know the status of all of the events and decisions that had arrows leading into the circle representing the event. For example, Figure 5.35 shows that initial U.S. and Soviet policy decision with respect to the Persian Gulf (decisions 2 and 3, and events 4 and 5) influenced subsequent events indirectly through their impact on the likelihood that there would be a significant increase in the level of revolutionary organization and activity in the area (event 7).

Defining the dependencies of the events and decisions in the Persian Gulf model, as shown in Figure 5.35, drastically reduced the number of probability assessments required of experts. In this case, the number of probability assessments was reduced from several million to approximately 100. Although experts disagreed over some of the influence linkages shown in Figure 5.35, none of the proposed changes would have resulted in a significantly different number of probability assessments. More importantly, the influence diagram in Figure 5.35 gave the experts a language with which to communicate their differences of opinion over the relationships between U.S. policy decisions and events that might have economic implications for the United States.



NOTE: See Table 5.1 for meaning of numbers.

FIGURE 5.35 AN APPLICATION OF AN INFLUENCE DIAGRAM

In addition to reducing the number of probability assessments required for the model, careful definitions of the dependencies between uncertain events greatly reduced the number of calculations required to manipulate the model. The influence diagram in Figure 5.35 pinpoints those portions of the decision tree in which redundant probability calculations occurred. By eliminating the redundancies, it was possible to reduce both the size of the decision tree and the number of calculations required to analyze the tree. When this was done, the model could be reduced from a decision tree with millions of decision and event nodes to a highly integrated decision tree with less than a thousand nodes.

Appendix A

PROBABILISTIC DEPENDENCE

## Appendix A

### PROBABILISTIC DEPENDENCE

To simplify the tasks of assessing and processing uncertain information it is often necessary to assume that random variables are independent. With this assumption, it is possible to assess the probability distribution for each random variable separately, and deal with relatively simple marginal probability distributions rather than complicated conditional distributions. Most large decision analysis projects contain within them some implicit or explicit assumptions about the independence of various random variables.

In order to describe independencies, we need to have a clear understanding of the types of independence that are possible. Two random variables are either dependent or independent. However, as we shall see, there are twenty-two different combinations of dependence and independence that can exist among three random variables. When there are more than three random variables, there are many different combinations of dependence and independence that exist among them.

This section describes the twenty-two combinations of dependence and independence that can exist among three random variables and gives an example of each combination. This topic has been discussed by Tribus [15]. However, Tribus considers only twelve of the twenty-two possible combinations of dependence and independence that can exist among three random variables.

#### 1. Independence Equation for Two Random Variables

When we assert that two random variables are independent, we are assuming that their joint probability distribution is equal to the product of the two marginal probability distributions. In other words, if

we say that random variables A and B are independent, then we are claiming that the following equation is true.

$$\langle AB \rangle = \langle A \rangle \langle B \rangle$$

Alternatively if we say that A and B are dependent, then we are claiming that the equation above is not true. To avoid writing an equation for each independence assertion, we will denote the assumption of independence between random variables A and B as follows:  $I(A,B)$ . If A and B are dependent, then we write  $\bar{I}(A,B)$ . If there are only two random variables, we must have  $I(A,B)$  or  $\bar{I}(A,B)$ .

## 2. Independence Equations for Three Random Variables

If we have three random variables-- A , B , and C --then there are ten possible independence equations. For example, we can assume that the three random variables are mutually independent, which means that their joint probability distribution is equal to the product of the three marginal distributions. Mutual independence is equivalent to the following equation.

$$\langle ABC \rangle = \langle A \rangle \langle B \rangle \langle C \rangle$$

If the three random variables are not mutually independent, then this equation is not true. In our notation mutual independence is written  $I(A,B,C)$ . If the three random variables are not mutually independent, we write  $\bar{I}(A,B,C)$ . The lack of mutual independence does not mean that the three random variables are mutually dependent, since other types of independence are possible. Mutual dependence means that there are no independencies among any of the random variables.

Another type of independence that can be asserted among the three random variables is that two of the three random variables are independent without regard to the third. When we have three random variables-- A , B , and C --we can assert  $I(A,B)$  or equivalently

$$\{AB\} = \{A\} \{B\}$$

This means that A and B are independent as long as we do not know C. Obviously, we can make similar independence assertions for A and C, or B and C.

Another type of independence assertion is that one of the random variables is independent of the distribution for the other two. If we assume that A is independent of B and C, then

$$\{ABC\} = \{A\} \{BC\}$$

In our notation, this independence equation is denoted  $I(A, BC)$ . This equation means that learning the value of A will not change the joint probability distribution for B and C. Alternatively, it means that learning B and C will not change the probability distribution for A. We could also assert that B is independent of A and C, or that C is independent of A and B.

When we learn the value of one of the three random variables, it changes our state of information and therefore can change our assumptions about independence for the remaining random variables. This allows us to make a different kind of independence assertion. For example, we can assume that when we know C, A and B are independent. The equation for this independence assertion is

$$\{AB|C\} = \{A|C\} \{B|C\}$$

When  $\{C\} \neq 0^*$ , the following equation is equivalent to the one above.

$$\{ABC\} = \frac{\{AC\} \{BC\}}{\{C\}}$$

---

\*This is not a restrictive condition since  $\{AB|C\}$ ,  $\{A|C\}$ , and  $\{B|C\}$  are not defined when  $\{C\} = 0$ .



This independence assumption is denoted  $I(A,B|C)$ . As before we can find some other independence assumptions by permuting the random variables. Thus, we can assume that B and C are independent when we know A, or we can assume that A and C are independent when we know B.

In this discussion it is assumed that independence can only be asserted among quantities that are each based on the same state of information. Independence among quantities that are based on different states of information is difficult to define since the different states of information may themselves affect the validity of the independence assumption. For example, can we assert independence between A by itself, and B given C? The independence equation corresponding to this statement is not well defined. We are trying to assert independence between  $\{A\}$  and  $\{B|C\}$ , and by analogy with the other independence equations we should be able to equate the product of these two quantities to some other quantity. However, it is not clear whether the product should equal  $\{AB\}$ ,  $\{AB|C\}$ , or something else.

In summary, there are ten possible independence equations or assertions that can be made for three random variables. These independence equations are listed below.

- (1)  $I(A,B,C) : \{ABC\} = \{A\} \{B\} \{C\}$
- (2)  $I(A,B) : \{AB\} = \{A\} \{B\}$
- (3)  $I(A,C) : \{AC\} = \{A\} \{C\}$
- (4)  $I(B,C) : \{BC\} = \{B\} \{C\}$
- (5)  $I(A,BC) : \{ABC\} = \{A\} \{BC\}$
- (6)  $I(B,AC) : \{ABC\} = \{B\} \{AC\}$
- (7)  $I(C,AB) : \{ABC\} = \{C\} \{AB\}$
- (8)  $I(A,B|C) : \{AB|C\} = \{A|C\} \{B|C\}$
- (9)  $I(A,C|B) : \{AC|B\} = \{A|B\} \{C|B\}$
- (10)  $I(B,C|A) : \{BC|A\} = \{B|A\} \{C|A\}$

### 3. Relationships Among Independence Equations

Since each of these independence equations can be either true or not true there would be  $2^{10}$  or 1,024 possible combinations of dependence and

independence among three random variables if it were not for the fact that some of the independence equations imply others. We can use the relationships among the ten independence equations to eliminate most of the 1,024 possible combinations, leaving twenty-two possible combinations of independence assertions that can exist among three random variables.

The first relationship among the independence equations is that mutual independence implies all of the other types of independence.

$$I(A,B,C) \rightarrow I(A,B), I(A,C), I(B,C), I(A,BC), I(B,AC), \\ I(C,AB), I(A,B|C), I(A,C|B), I(B,C|A)$$

To prove this relationship we start with the equation for mutual independence:  $\{ABC\} = \{A\} \{B\} \{C\}$ . If we integrate both sides of this equation over all possible values of  $C$ , we have  $\{AB\} = \{A\} \{B\}$ . This proves that  $A$  and  $B$  are independent without regard to  $C$ , or  $I(A,B)$ . Similarly, by integrating over all possible values of  $B$  and  $A$ , we find that  $\{AC\} = \{A\} \{C\}$  and  $\{BC\} = \{B\} \{C\}$ . These two equations are equivalent to  $I(A,C)$  and  $I(B,C)$ . If we substitute these equations into the equation for mutual independence, we have  $\{ABC\} = \{A\} \{BC\} = \{B\} \{AC\} = \{C\} \{AB\}$ , or  $I(A,BC)$ ,  $I(B,AC)$ , and  $I(C,AB)$ . By using these results we show that

$$\{AB|C\} = \frac{\{ABC\}}{\{C\}} = \frac{\{A\} \{B\} \{C\}}{\{C\}} = \{A\} \{B\} \\ = \{A|C\} \{B|C\} \text{ when } \{C\} \neq 0^*$$

This is equivalent to  $I(A,B|C)$ . In exactly the same way we can prove  $I(A,C|B)$  and  $I(B,C|A)$  when  $B \neq 0$  and  $A \neq 0$ . Thus, mutual independence implies all other types of independence.

---

\*This is not a restrictive condition since  $\{AB|C\}$ ,  $\{A|C\}$ , and  $\{B|C\}$  are not defined when  $\{C\} = 0$ .

The other relationships among the independence equations are listed below. The proofs for these relationships are similar to the one for mutual independence, and they are outlined briefly.

1.  $I(A,B,C) - I(A,B), I(A,C), I(B,C), I(A,BC), I(B,AC), I(C,AB), I(A,B|C), I(A,C|B), I(B,C|A)$
2.  $I(A,BC) - I(A,B), I(A,C)$   
 Proof:  $I(A,BC) - \{ABC\} = \{A\} \{BC\} = \frac{\{ABC\}}{C} = \frac{\{A\} \{BC\}}{C}$   
 $= \{AB\} = \{A\} \{B\} = I(A,B)$
3.  $I(A,BC) - I(A,B|C), I(A,C|B)$   
 Proof:  $I(A,BC) - I(A,B), I(A,C)$   
 $= \{AB\} = \{A\} \{B\}, \{AC\} = \{A\} \{C\}$   
 $\{A\} = \frac{\{AB\}}{\{B\}} = \frac{\{AC\}}{\{C\}}$   
 $I(A,BC) - \{ABC\} = \{A\} \{BC\} = \frac{\{AB\} \{BC\}}{\{B\}} = \frac{\{AC\} \{BC\}}{\{C\}}$   
 $= I(A,B|C), I(A,C|B)$
4.  $I(A,BC), I(B,C) - I(A,B,C)$   
 Proof:  $I(B,C) - \{BC\} = \{B\} \{C\}$   
 $I(A,BC) - \{ABC\} = \{A\} \{BC\} = \{A\} \{B\} \{C\}$   
 $= I(A,B,C)$
5.  $I(A,BC), I(B,AC) - I(A,B,C)$   
 Proof:  $I(B,AC) - I(B,C)$   
 $I(A,BC), I(B,C) - I(A,B,C)$
6.  $I(A,BC), I(B,C|A) - I(A,B,C)$   
 Proof:  $I(A,BC) - I(A,B), I(A,C)$   
 $= \{AB\} = \{A\} \{B\}, \{AC\} = \{A\} \{C\}$   
 $I(B,C|A) - \{ABC\} = \frac{\{AB\} \{AC\}}{\{C\}} = \{A\} \{B\} \{C\}$   
 $= I(A,B,C)$
7.  $I(B,C|A), I(A,B) - I(B,AC)$   
 Proof:  $I(A,B) - \{AB\} = \{A\} \{B\}$   
 $I(B,C|A) - \{ABC\} = \frac{\{AB\} \{AC\}}{\{A\}} = \{B\} \{AC\}$   
 $= I(B,AC)$

#### 4. Combinations of Independence Equations

The twenty-two possible combinations of independence equations are shown graphically in Figure A.1. The remaining 1,002 combinations of

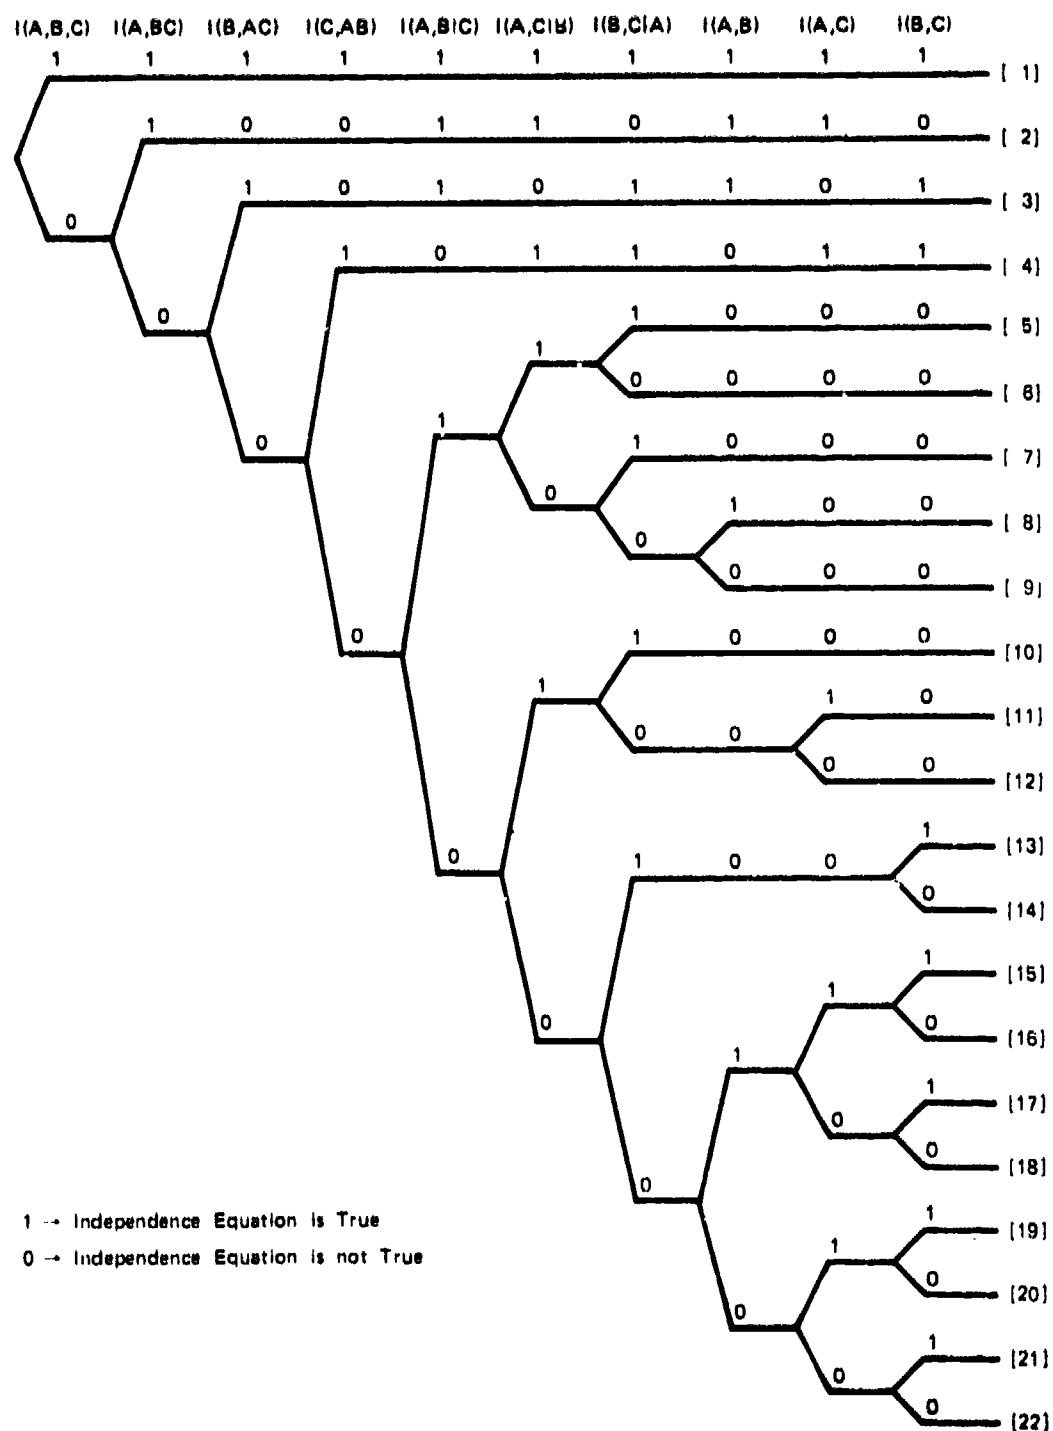


FIGURE A.1 COMBINATIONS OF INDEPENDENCE EQUATIONS

independence equations are not possible since they violate one of the relationships above. The twenty-two combinations of independence equations have been labeled with numbers in square brackets at the right of A.1. The first combination corresponds to mutual independence and the twenty-second combination corresponds to mutual dependence. The other twenty combinations correspond to various intermediate levels of independence among three random variables.

The combinations of independence equations are demonstrated below with simple examples based on flipping four types of coins: fair, biased, thick and magnetized coins. The random variables in each example describe whether a head or tail results from flipping the coins. These random variables are discrete, but it is easy to generalize the examples to include continuous random variables. Examples are provided for only 10 of the 22 possible combinations of independence equations, but any combination can be demonstrated by permuting the random variables in these examples.

Example (1): Combination [1] in Figure A.1

$I(A,B,C)$ ,  $I(A,BC)$ ,  $I(B,AC)$ ,  $I(C,AB)$ ,  $I(A,B|C)$ ,  
 $I(A,C|B)$ ,  $I(B,C|A)$ ,  $I(A,B)$ ,  $I(A,C)$ ,  $I(B,C)$

In this example A , B , C correspond to the outcomes of flipping three fair coins. The joint probability mass function for A , B , and C is:

		B		B			
		H	T	H	T		
A	H	0.125	0.125	0.125	0.125	H	A
	T	0.125	0.125	0.125	0.125	T	
		C = H		C = T			

Example (2): Combination [2] in Figure A.1

$\bar{I}(A,B,C)$ ,  $I(A,BC)$ ,  $\bar{I}(B,AC)$ ,  $\bar{I}(C,AB)$ ,  $I(A,B|C)$ ,  
 $I(A,C|B)$ ,  $\bar{I}(B,C|A)$ ,  $I(A,B)$ ,  $I(A,C)$ ,  $\bar{I}(B,C)$

In this example A corresponds to the outcome of flipping a fair coin, and B and C correspond to the outcomes of flipping two magnetic coins. The probability that a head will occur when the first magnetic coin is tossed is 50%. However, there is a 60% chance that the second magnetic coin will land with the same side up as the first magnetic coin. Neither of the magnetic coins are affected by the outcome of tossing the fair coin. The conditional probabilities for this example are:

$$\{A=H\} = \{B=H\} = 0.5$$

$$\{C=H|B=H\} = \{C=T|B=T\} = 0.6$$

The joint probability mass function for the three random variables is:

C				C			
		H	T			H	T
B	H	0.15	0.10	0.15	0.10	H	B
	T	0.10	0.15	0.10	0.15	T	
A = H				A = T			

Example (3): Combination [5] in Figure A.1

$\bar{I}(A,B,C)$ ,  $\bar{I}(A,BC)$ ,  $\bar{I}(B,AC)$ ,  $\bar{I}(C,AB)$ ,  $I(A,B|C)$ ,  
 $I(A,C|B)$ ,  $I(B,C|A)$ ,  $\bar{I}(A,B)$ ,  $\bar{I}(A,C)$ ,  $\bar{I}(B,C)$

In this example A, B, and C each correspond to the same flip of a fair coin. The conditional probabilities for this example are:

$$\{A=H\} = 0.5$$

$$\{B=H|A=H\} = \{B=T|A=T\} = 1.0$$

$$\{C=H|A=H\} = \{C=T|A=T\} = 1.0$$

The joint probability mass function for the three random variables is:

B				B			
		H	T			H	T
A	H	0.5	0			0	0
	T	0	0			0	0.5
C = H				C = T			

Example (4): Combination [6] in Figure A.1

$\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), I(A,B|C),$   
 $I(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$

In this example A corresponds to the outcome of flipping a fair coin. If A is a head, then B and C must also be heads. If A is a tail, then B and C both correspond to the same flip of another fair coin. The conditional probabilities for this example are:

$$\begin{aligned} \{A=H\} &= 0.5 \\ \{B=H|A=H\} &= 1.0, \quad \{B=T|A=T\} = 0.5 \\ \{C=H|B=H\} &= \{C=T|B=T\} = 1.0 \end{aligned}$$

The joint probability mass function for the three random variables is:

C				C					
		H	T			H	T		
B	H	0.5	0			0.25	0	H	B
	T	0	0			0	0.25	T	
A = H				A = T					

Example (5): Combination [8] in Figure A.1

$\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), I(A,B|C),$   
 $\bar{I}(A,C|B), \bar{I}(B,C|A), I(A,B), \bar{I}(A,C), \bar{I}(B,C)$

In this example C corresponds to the outcome of flipping a thick coin (a cylinder) that can land on its edge in addition to heads or tails. Possible outcomes for C are heads (H), tails (T), and edge (E). A and

B correspond to the outcome of flipping two biased coins, where the bias of each coin depends on the outcome of C. The conditional probabilities for this example are:

$$\begin{aligned} \{A=H|C=H\} &= 0.4, \quad \{B=H|C=H\} = 0.6 \\ \{A=H|C=T\} &= 0.5, \quad \{B=H|C=T\} = 0.1 \\ \{A=H|C=E\} &= 0.7, \quad \{B=H|C=E\} = 0.6 \\ \{C=H\} &= \{C=T\} = 0.4, \quad \{C=E\} = 0.2 \end{aligned}$$

The joint probability mass function for the three random variables is:

		B		B		B			
		H	T	H	T	H	T		
A	H	0.096	0.064	0.020	0.180	0.084	0.056	H	A
	T	0.144	0.096	0.020	0.180	0.036	0.024	T	
		C = H		C = T		C = E			

It can be shown that this combination of independence assertions cannot exist for three binary events. Thus it was necessary for C to have three possible outcomes in the example above.

Example (6): Combination [9] in Figure A.1

$$\begin{aligned} &\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), I(A,B|C), \\ &\bar{I}(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C) \end{aligned}$$

In this example C corresponds to the outcome of flipping a fair coin. A and B correspond to the outcomes of flipping two biased coins, where the direction in which the coins are biased depends on C. The conditional probabilities for this example are:

$$\begin{aligned} \{A=H|C=H\} &= \{B=H|C=H\} = 0.6 \\ \{A=H|C=T\} &= \{B=H|C=T\} = 0.4 \\ \{C=H\} &= 0.5 \end{aligned}$$



The joint probability mass function for the three random variables is

		B		B			
		H	T	H	T		
A	H	0.18	0.12	0.08	0.12	H	A
	T	0.12	0.08	0.12	0.18		T
		C = H		C = T			

Example (7): Combination [15] in Figure A.1

$\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$   
 $\bar{I}(A,C|B), \bar{I}(B,C|A), I(A,B), I(A,C), I(B,C)$

In this example C corresponds to the outcome of flipping a fair coin. A and B correspond to the outcomes of flipping two magnetized coins, where the direction in which B is magnetized depends on C. The conditional probabilities for this example are:

$$\{C=H\} = \{B=H\} = 0.05$$

$$\{A=H|B=H, C=H\} = \{A=T|B=T, C=H\} = 0.6$$

$$\{A=H|B=H, C=T\} = \{A=T|B=T, C=T\} = 0.4$$

The joint probability mass function for the three random variables is:

		B		B			
		H	T	H	T		
A	H	0.15	0.10	0.10	0.15	H	A
	T	0.10	0.15	0.15	0.10		T
		C = H		C = T			

Example (8): Combination [16] in Figure A.1

$\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$   
 $\bar{I}(A,C|B), \bar{I}(B,C|A), I(A,B), I(A,C), \bar{I}(B,C)$

In this example A corresponds to the outcome of flipping a fair coin. B and C correspond to the outcomes of flipping two magnetized coins,

where the outcome of A is used to determine how strongly the coins are magnetized. The conditional probabilities for this example are:

$$\{A=H\} = \{B=H\} = 0.5$$

$$\{C=H|B=H, A=H\} = \{C=T|B=T, A=H\} = 0.8$$

$$\{C=H|B=H, A=T\} = \{C=T|B=T, A=T\} = 0.6$$

The joint probability mass function for the three random variables is:

		C		C			
		H	T	H	T		
B	H	0.20	0.05	0.15	0.10	H	B
	T	0.05	0.20	0.10	0.15		T
		A = H		A = T			

Example (9): Combination [18] in Figure A.1

$\bar{I}(A, B, C), \bar{I}(A, BC), \bar{I}(B, AC), \bar{I}(C, AB), \bar{I}(A, B|C),$   
 $\bar{I}(A, C|B), \bar{I}(B, C|A), I(A, B), \bar{I}(A, C), \bar{I}(B, C)$

In this example A and B correspond to outcomes of flipping two fair coins. C corresponds to the outcome of flipping a biased coin where the amount that the coin is biased depends on A and B. The conditional probabilities for this example are:

$$\{A=H\} = \{B=H\} = 0.5$$

$$\{C=H|A=H, B=T\} = 0.4$$

$$\{C=H|A=H, B=H\} = \{C=H|A=T, B=H\} = 0.6$$

$$\{C=H|A=H, B=H\} = 0.8$$

The joint probability mass function for the three random variables is:

		B		B			
		H	T	H	T		
A	H	0.20	0.15	0.05	0.10	H	A
	T	0.15	0.10	0.10	0.15		T
		C = H		C = T			

Example (10): Combination [22] in Figure A.1

$$\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C), \\ \bar{I}(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$$

In this example A , B , and C , correspond to the outcome of flipping three magnetic coins. There is a 50% chance that the first coin flipped will come up heads, and there is a 60% chance that the second coin flipped will land with the same side up as the first. If the first two coins are both heads or both tails, there is an 80% chance that the third coin flipped will have the same outcome; otherwise the probability of heads on the third flip is 50%. The conditional probabilities for this example are:

$$\{A=H\} = 0.5$$

$$\{B=H|A=H\} = \{B=T|A=T\} = 0.6$$

$$\{C=H|A=H,B=H\} = \{C=T|A=T,B=T\} = 0.8$$

$$\{C=H|A=H,B=T\} = \{C=H|A=T,B=H\} = 0.5$$

The joint probability mass function for the three random variables is:

B				B			
		H	T			H	T
A	H	0.24	0.10	0.06	0.10	H	A
	T	0.10	0.06	0.10	0.24	T	
C = H				C = T			

#### E. An Additional Relationship Among the Independence Equations

Some of the combinations of independence equations require the joint distribution for the three random variables to contain several zeros. This means that certain combinations of outcomes for the three random variables are not possible even though each of the random variables can individually assume the same outcomes. In Example (3), A , B , and C can each be heads or tails, but it is not possible for one to be a head when another is a tail.

If we assume that the joint distribution cannot equal zero when the individual marginal distributions are not zero, it is possible to prove an additional relationship among the ten independence equations.

$$I(A,B|C), I(A,C|B) - I(A,BC) \text{ if } \{BC\} \neq 0$$

The assumption the  $\{BC\}$  is not zero allows us to divide by this quantity in the following proof:

$$\begin{aligned} I(A,B|C), I(A,C|B) - \{ABC\} &= \frac{\{AC\} \{BC\}}{\{C\}} - \frac{\{AB\} \{BC\}}{\{B\}} \\ &= \{AC\} \{B\} - \{AB\} \{C\} = \{AC\} \{B\} - \{AB\} \{C\} \\ &= \{A\} \{B\} - \{AB\} = I(A,B) \\ I(A,C|B), I(A,B) &= I(A,BC) \end{aligned}$$

If this relationship is added to the seven discussed above, it can be used to eliminate four of the twenty-two possible combinations of independence equations. The four combinations that are eliminated are numbered [5], [6], [7], and [10] in Figure A.1. The remaining eighteen combinations of independence equations are still possible.

#### F. Encoding the Twenty-Two Combinations of Independence Equations

A subjects' state of information about several uncertain quantities can be represented by a wide variety of possible combinations of independence equations, even when the problem contains as few as three random variables. One of the principal motivations for assuming that random variables are independent is to limit the amount of probability encoding required to specify the joint probability distribution for all random variables. Although there are many possible combinations of independence equations, the degree of difficulty associated with assessing the uncertainties necessary to specify the joint distribution can be determined by some very simple properties of the independence equations that are assumed to be true. We can place the ten possible independence equations for three random variables in the four categories as shown below:

Category 1:  $I(A,B,C) : \{ABC\} = \{A\} \{B\} \{C\}$

Category 2:  $I(A,BC) : \{ABC\} = \{A\} \{BC\}$   
 $I(B,AC) : \{ABC\} = \{B\} \{AC\}$   
 $I(C,AB) : \{ABC\} = \{C\} \{AB\}$

Category 3:  $I(A,B|C) : \{AB|C\} = \{A|C\} \{B|C\}$   
 $I(A,C|B) : \{AC|B\} = \{A|B\} \{C|B\}$   
 $I(B,C|A) : \{BC|A\} = \{B|A\} \{C|A\}$

Category 4:  $I(A,B) : \{AB\} = \{A\} \{B\}$   
 $I(A,C) : \{AC\} = \{A\} \{C\}$   
 $I(B,C) : \{BC\} = \{B\} \{C\}$

The degree of difficulty associated with assessing the probabilities needed to specify the joint distribution for  $A$ ,  $B$ , and  $C$  depends only on which categories of independence assumptions contain equations that are assumed to be true. As a measure of the degree of difficulty associated with the assessment problem, we will assume that  $A$ ,  $B$ , and  $C$  are each discrete random variables with probability mass functions that contain  $n$  possible outcomes; we will then determine the minimum number of probabilities required to specify the joint mass function.

If the three random variables are assumed to be mutually independent (Category 1), then the joint distribution can be determined by assessing the three marginal distributions and multiplying them together. To specify each of the marginal probability mass functions, we would need to assess  $n$  probabilities. Therefore, to determine the joint probability mass function for  $A$ ,  $B$ , and  $C$ , we would need to assess  $3n$  probabilities.

If our state of information about the three-random variables is such that we can not assume mutual independence, but can assume that one of the independence equations in Category 2 is true, we can determine the joint probability density function by assessing one of the marginal distributions and the joint distribution for the two remaining random variables. We need to assess  $n$  probabilities to determine the marginal distribution, and  $n^2$  probabilities to determine the joint distribution for two random variables. Thus, when we can assume that one of the independence

equations in Category 2 is true, but that the independence equation in Category 1 is not true, we would need to assess  $(n^2 + n)$  probabilities. This situation occurs for the combinations of independence equations numbered [2], [3], and [4] in Figure A.1.

If our state of information about the three random variables is such that we cannot assume that the independence equations in Categories 1 and 2 are true, but one of the independence equations in Category 3 is true, we can determine the joint probability density function by assessing the marginal distribution for one of the three random variables and the conditional distributions for the other two random variables, given the first. In this situation, we would need to assess  $n$  probabilities for the marginal distribution, and  $n^2$  probabilities for each of the two conditional distributions. Thus, when the independence equations in Categories 1 and 2 are not true, but one of the independence equations in Category 3 is true, we will need to assess  $(2n^2 + n)$  probabilities. This situation occurs for the combinations of independence equations numbered [5] through [14] in Figure A.1.

If our state of information about the three random variables is such that none of the independence equations in Categories 1, 2, and 3 are true, but one of the independence equations in Category 4 is true, then we can determine the joint density function by assessing two marginal distributions and the conditional distribution for the third random variable, given the first two. However, to assess the conditional distribution, we would need to assess  $n^3$  probabilities. Since we can assess the joint distribution for all three random variables with  $n^3$  probabilities, we can minimize the number of probabilities assessed by doing so. We could also assess the joint distribution for all three random variables by assessing the  $n^3$  probabilities in the case where none of the independent equations are true. Thus, when none of the independence equations in Categories 1, 2, and 3 are true, we will need to assess at least  $n^3$  probabilities.

The number of probabilities that must be assessed to determine the joint distribution for all three random variables is summarized in the following table as a function of the categories of independence equations.

Categories of Independence Equations	Number of Probabilities Assessed
--	--

1	$3n$
2	$n^2 + n$
3	$2n^2 + n$
4	$n^3$

The number of probabilities that must be assessed to determine the joint distribution is a rough measure of the degree of difficulty associated with the encoding process. For a three-variable problem where each variable has many possible values--corresponding to large values of  $n$ --mutual independence (Category 1) or partial independence (Categories 2 and 3) can be powerful simplifying assumptions. Where more than three variables are involved the independence assumptions are ever more useful.

The number of probabilities that must be assessed to determine the joint distribution is a rough measure of the degree of difficulty associated with the encoding process. For a three-variable problem where each variable has many possible values--corresponding to large values of  $n$ --mutual independence (Category 1) or partial independence (Categories 2 and 3) can be powerful simplifying assumptions. Where more than three variables are involved the independence assumptions are even more useful.

Appendix B

APPROXIMATE METHODS OF CALCULATING THE VALUE OF INFORMATION



## Appendix B

### APPROXIMATE METHODS OF CALCULATING THE VALUE OF INFORMATION

The purpose of this section is to explain how approximate methods can be used to determine the value of information. The approximate value-of-information calculations are based on deterministic sensitivities, and decisions to increase or decrease the size of a model can be based on the approximate values of information.

#### 1. Introduction

Decision problems may be grouped according to the the number of stages they contain. We define a half-stage problem as one without decision variables. A single-stage problem has any number of decision variables followed by any number of state variables. Multistage problems are characterized by decision variables separated by state variables. Figure B.1 classified typical problems. Any decision problem can be solved to any desired accuracy using trees; however, in many problems certain "smoothness" properties can be exploited to find answers as accurate as those from trees at a fraction of the computational cost. For one-half or one-stage problems specific procedures can be programmed that are adequate for the majority of problems. Two-stage procedures exist for a class of problems. Practical procedures do not exist for many-stage problems

For simplicity we will treat the case of risk indifference. However, Rice [14] has shown that only minor modifications are required to extend the framework to include risk aversion.

A value function and a deterministic model, can be expanded in a Taylor series about the mean of the state variables and the corresponding optimum setting of the decision variables. This approach is exact for a single-stage problem that is quadratic in continuous state and decision

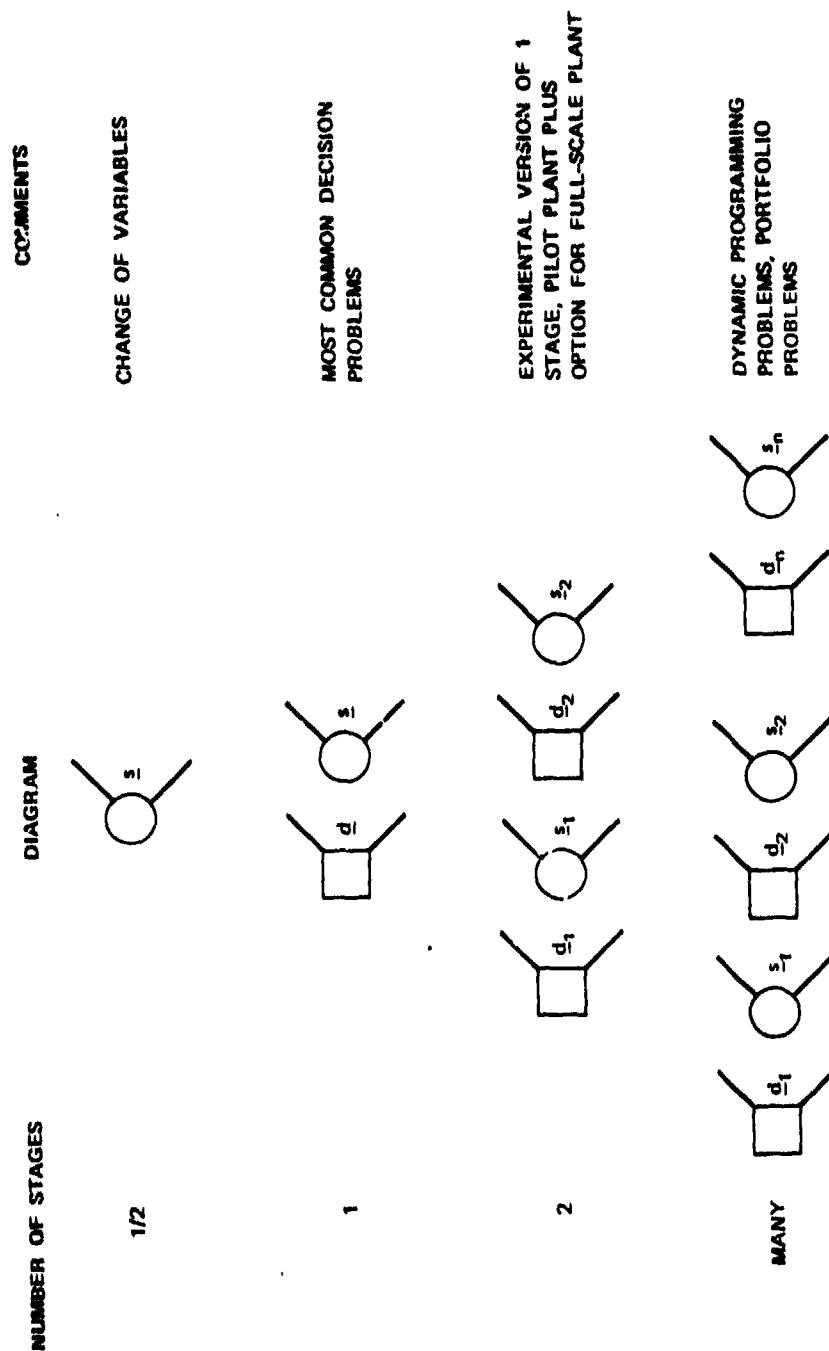


FIGURE B.1 CLASSIFICATION OF DECISION PROBLEMS

variables and for discrete decision problems where the value function is a linear function of the state variables.

Howard has developed and documented this approach for change of variable problems in his course notes for EES 221 [7] at Stanford University, and for one-stage problems in his paper Proximal Decision Analysis [6]. His results are algebraic expressions for the mean and variance of the profit lottery and for the value of clairvoyance on the state variables. Means, variances, and covariances for the state variables, as well as the partial derivatives of profit with respect to the state and decision variables are required inputs. To automate Howard's methodology, a computer would be given the profit function and the moments of the state variables. The computer would automatically run sensitivities to estimate the required partial derivatives.

Rice has generalized Howard's approach. The general approach also starts with sensitivities and ends with the value of clairvoyance on each state variable. However, the middle step is eliminated, going directly from sensitivities to value of clairvoyance without evaluating partial derivatives. For continuous problems, ones with continuous state and decision variables, the two methods are equivalent. The advantage of the direct method is that it also works for problems with discrete decision and/or state variables and with discontinuous value functions.

## 2. Computerizing the Single-Stage Model

For a computer to find the value of clairvoyance for a state variable given only a single-stage profit function and the moments (or distributions) of the state variables, two conditions must hold:

- (i) The state variables must be probabilistically independent of decision variables.
- (ii) The value structure must be of the form  $v(s,d)$ : a deterministic model which assigns a single profit measure to each complete vector of state and decision variables.

Neither condition is restrictive. The Entrepreneur's Problem in Proximal Decision Analysis [6] demonstrates how apparent violations of (i) can be

rectified by reformulating the deterministic model. The model form  $v(s,d)$  in (ii) is common to many discrete as well as continuous problems.

The output for the computer program is the approximate value of information. The expected value of clairvoyance on the  $i$ 'th state variable is the prior expectation of the conditional expected value given  $s_i$  less the prior expected value:

$$\langle v_{c_i} | E \rangle = \langle \langle v | s_i, E \rangle \rangle - \langle v | E \rangle$$

Alternatively, we can focus on changes of decision by subtracting the prior expected value from the conditional expected value before taking the second expectation:

$$\langle v_{c_i} | E \rangle = \langle \langle v | s_i, E \rangle - \langle v | E \rangle | E \rangle$$

The value of clairvoyance is the expected increase in value from making decisions after  $s_i$  is revealed rather than before. We define the quantity:

$$\langle v | s_i, E \rangle - \langle v | E \rangle$$

as the stochastic compensation: "stochastic" because each term is an expected value, and "compensation" because the decision is reoptimized to compensate for the departure of the  $i$ 'th state variable from its mean. Using the new terminology, the value of clairvoyance on the  $i$ 'th state variable is the expected stochastic compensation.

To approximate the value of clairvoyance, deterministic compensation may be substituted for stochastic compensation.

### 3. Steps to Compute the Approximate Value of Information

To calculate the approximate value of information for the single stage problem, a computer must complete the following steps:

- 1) Accept the problem specification.
- 2) Solve for the deterministic optimum decision.
- 3) Perform deterministic open loop sensitivities.
- 4) Perform deterministic closed loop sensitivities.
- 5) Generate compensation functions.
- 6) Compute the expected compensation.

The procedure is remarkably robust. It works regardless of whether the state and decision variables are continuous or discrete. However, because of difference in optimization techniques the algorithms to achieve step 1-6 are different for discrete and continuous decision variables.

#### 4. Discrete Decision Variable

For a discrete decision variable steps 2-4 can be performed simultaneously. As illustrated in Figure B.2 for two decision alternatives, the value  $v$  must be computed for each alternative at  $\bar{s}$ , and at  $\bar{s} \pm s$ . Then the open-loop sensitivity is the curve with the highest value at  $\bar{s}$ , the one for  $d_1$  in the example. The closed loop sensitivity is the upper boundary of the curves. In Figure B.2 it is the  $d_1$  open-loop sensitivity to the right of the crossover point and the  $d_2$  open-loop sensitivity to the left of the crossover point. The compensation is the difference between the closed and open loop sensitivities, plotted in Figure B.3 for the example. The expected value of clairvoyance is found by numerically integrating the product of the compensation plot and the distribution of the state variable.

The extension of the procedure to many discrete alternatives is straightforward. There is one open-loop sensitivity for each alternative. The closed loop sensitivity is the concave envelope of the open-loop sensitivities. The compensation is the difference between the closed-loop and open-loop sensitivities as before.

#### 5. Continuous Decision Variables

For continuous decision variables or for discrete decision variables with many alternatives, the computer will generate conventional

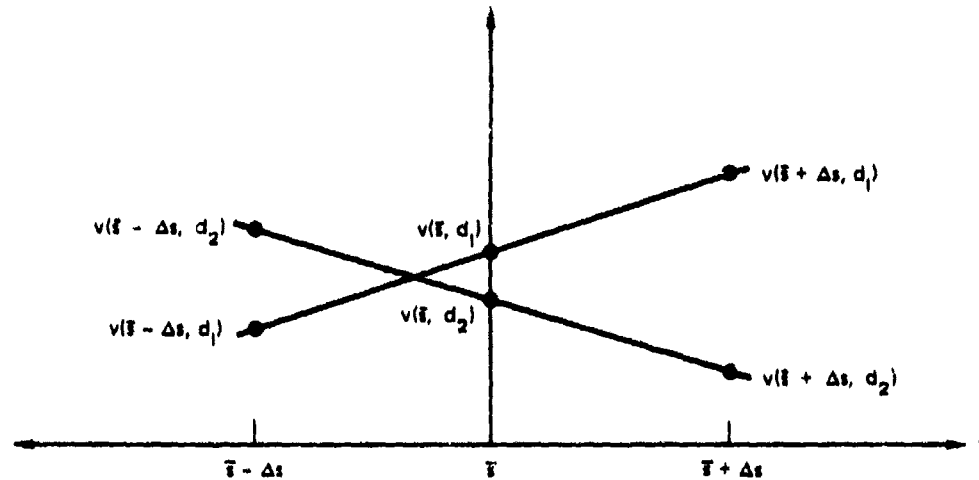


FIGURE B.2 OPEN AND CLOSED LOOP SENSITIVITIES

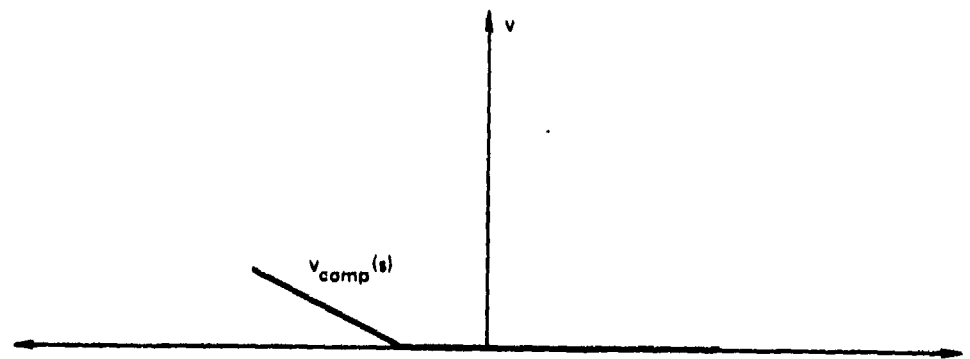
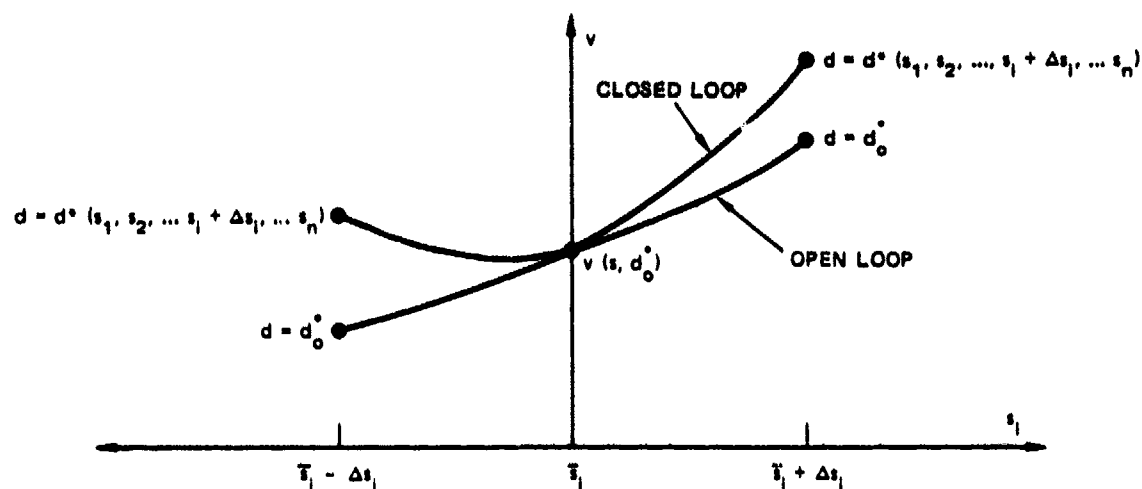
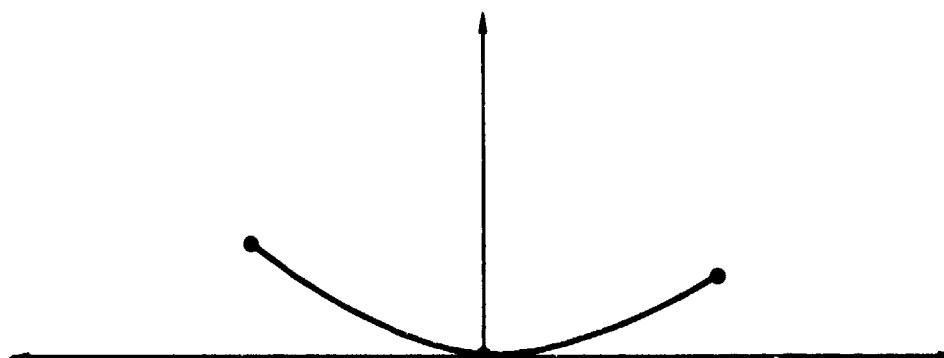


FIGURE B.3 COMPENSATION FOR A DISCRETE DECISION VARIABLE

sensitivities. Starting with the base case where all state variables are set to their means and all decision variables are set at the deterministic optimum, the computer successively adds and subtracts an increment to the  $i$ 'th state variable. As illustrated in Figure B.4a, the decision is reoptimized at each point for the closed loop sensitivity and held at the deterministic optimum for the open loop sensitivity. The compensation plotted in Figure B.4b is the difference between the open-loop and closed-loop sensitivities. To compute the expected compensation the computer will approximate the compensation with a quadratic or other curve form and perform numerical integration.



(a) OPEN AND CLOSED LOOP SENSITIVITIES

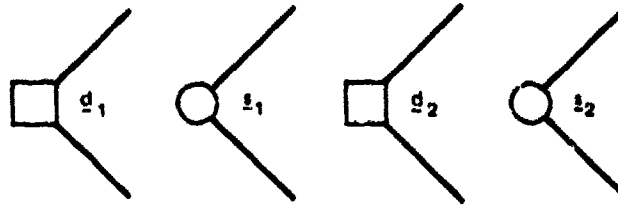


(b) COMPENSATION

FIGURE B.4 COMPUTATION FOR CONTINUOUS DECISION VARIABLES

#### F. Two Stage Problems

Merkhofer [10] has shown that under certain conditions the two-stage continuous decision problem illustrated in Figure B.5 is solvable by these methods. As discussed in The Economics of Decision Making [14], the only important terms in the value function are the G matrix of second partial derivatives of value with respect to the  $i$ 'th state variable and the  $j$ 'th decision variable and the H matrix of second partial derivatives of value with respect to decision variable  $i$  and decision variable  $j$ . These matrices may be positioned as shown in Figure B.5. The value of compensation for this model is defined to be the difference between the open-loop



NOTE:  $\{s_2 | s_1, d_1, E\} = \{s_2 | E\}; \{s_1 | d_1, E\} = \{s_1 | E\}$

(a) THE SCHEMATIC TREE

$$\begin{aligned}
 v(2, d) &= v(s_1, s_2, d_1, d_2) \\
 &= s' G d + d' H d \\
 &= \begin{bmatrix} s_1 & s_2 \end{bmatrix} \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + \begin{bmatrix} d_1 & d_2 \end{bmatrix} \begin{bmatrix} H_{11} & H_{12} \\ H_{12} & H_{22} \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix}
 \end{aligned}$$

(b) VALUE FUNCTION

FIGURE D.5 THE TWO STAGE QUADRATIC DECISION PROBLEM

sensitivity of the value function to the state variables  $s_1$  and the partially closed loop sensitivity in which the decision variables  $d_1$  are held at their deterministic optimums but the decision variables  $d_2$  are continuously optimized as  $s_1$  is varied. The expected deterministic compensation is calculated to be:

$$\langle v_{\text{comp}} | E \rangle = -(1/2) \langle s_1' G_{12} H_{22}^{-1} G_{12}' s_1 | E \rangle$$

The expected value of the optimal decision strategy will equal the expected deterministic compensation under either of the following conditions:

1.  $s_2$  is probabilistically independent of  $s_1$ .
2.  $G_{22}$  is a zero matrix and  $\langle s_2 | s_1, e \rangle | e \rangle = 0$ .

(Non-observable state variables are deterministically independent of the flexible decision variables and the prior expectation is that the posterior mean of  $s_2$  will not be shifted by knowledge of  $s_1$ .)

These conditions are frequently satisfied in practice.



Appendix C

THE VALUE OF DECISION-DEPENDENT INFORMATION

## Appendix C

### THE VALUE OF DECISION-DEPENDENT INFORMATION

The idea of valuing perfect information has appeared in many treatments of decision making under uncertainty. Most often the example being treated represents a simple hypothetical situation. The informational structure that is being captured in probability assignments is straightforward and the assumptions regarding the probabilistic structure, such as types of independence among both controlled and uncontrolled variables, are implicit in the problem statement. However, the correct computation of the value of information can be elusive on both conceptual and numerical bases. The concept of clairvoyance will lead us to the construction of detailed information models and to the exploration of their precise interpretation and use.

#### 1. The Primary Decision

For illustration, let us play the role of analysts for a space mission designed to land a remotely controlled experimental apparatus on the surface of Mars. We have thoroughly analyzed the mission and have summarized our total state of information by assigning a 0.6 probability that the Mars mission will be successful and a corresponding 0.4 probability that it will be a failure. Also we have analyzed the values to be derived from the mission and have put them in monetary units--millions-of-dollars, for example. Let's assume that the value of a successful Mars mission is 50 units and the value of an unsuccessful one is 10 units. A positive value might be attributed to a failure because attempting the mission has important social value and even a failure will provide knowledge for a better design on the next attempt.

Unexpectedly, several months before launch another nation announces that it will attempt a similar mission to Venus in about one year. Because

of the competitive nature of the space race and the important foreign policy implications of technological leadership, we realize that the value of changing our destination and successfully landing on Venus would be quite high. On the other hand, if we attempt to land on Venus and fail we would look foolish for diverting the program, and in any case we would set back the timetable for our extensive Martian exploration program at least two years. When all of these factors are evaluated we find that a successful landing on Venus is worth 100 units and a failure costs 10 units. To our surprise, when we check the feasibility of diverting the mission we find that because of modular design only a few important, but thoroughly tested, components of the landing system need to be changed, and the mission engineers assign a 0.6 probability of success regardless of destination.

From these assessments, we can lay out the primary decision tree of Figure C.1. Along each outcome branch emanating from a chance node we have written the conditional probability of following that path, and near each node we have written the value, either assigned or derived, of being at the point in the program represented by that node. We see that the expected value of going to Mars is 34, while the expected value of going to Venus is 56. Thus, in order to maximize the expected value of the mission we decide to go to Venus.

## 2. Value of Perfect Information

We might wish to use the decision tree to investigate the possibility of gathering new information before we make the final decision. To do this we can use the value of perfect information as an upper bound for the value of less complete information gathering programs. Most analysts, when presented with Figure C.1 and asked to derive the value of perfect information, reverse the order of decision and chance nodes in Figure C.1 to produce the tree shown in Figure C.2. With the latter tree we learn first whether the mission will succeed or fail, and then we decide on the destination. If we know the mission will succeed, we send it to Venus for 100 units of value, and if it will fail, we send it to Mars for 10 units of value. Using the original probability of success (0.6)

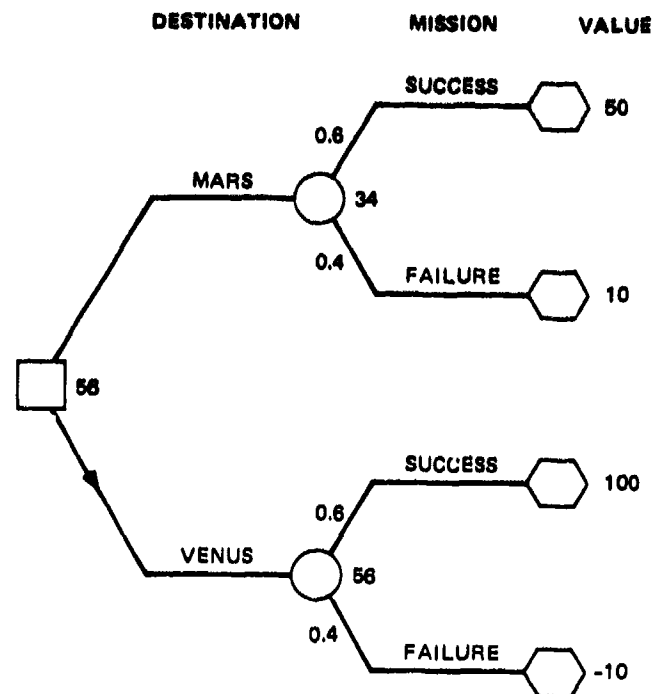


FIGURE C.1 THE PRIMARY DECISION TREE

as the probability that the information will predict a success, we obtain an expected value of 64 units with perfect information. Subtracting 56 units for the value of the primary decision problem, we obtain 8 units for the value of perfect information.

What might be wrong with this approach? Suppose that perfect information revealed that the mission would succeed on Mars but fail on Venus, or vice-versa. These possibilities do not appear in Figure C.2. To correct this omission we might draw a new tree for the value with perfect information as illustrated in Figure C.3. First we learn one of four possible predictions consisting of the four combinations of success or failure on Mars and Venus. Then we make the best decisions given this information, as indicated in the decision tree.

In order to assign probabilities one might reason that, since landings on Mars or Venus appear in separate portions of the primary tree of Figure C.1, the events must be independent and the probabilities should be multiplied as shown in Figure C.3. This will yield a value of 73.6

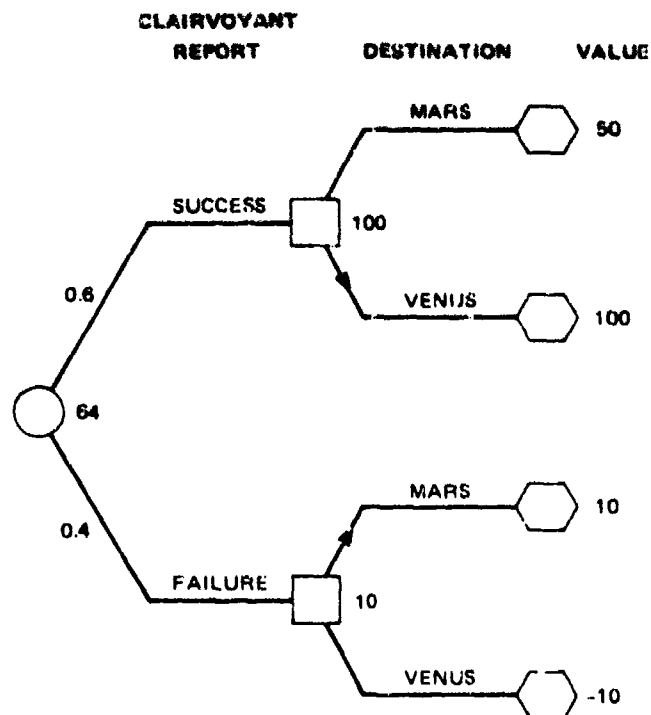


FIGURE C.2 TYPICAL DECISION TREE FOR DETERMINING THE VALUE OF PERFECT INFORMATION

with perfect information, and subtracting the 56 unit value of the primary decision it yields a 17.6 value of perfect information.

However, the independence assumption must be questioned. Since we can only send the mission to a single destination, might the events be mutually exclusive? If we simply try to assign the probabilities directly we are tempted to phrase confusing questions like "What is the probability we will succeed on both Mars and Venus?" or "If we learned we had failed on Mars what probability would we assign to success on Venus?" The trouble stems from the fact that we have only one rocket and it is difficult to consider sending it to both destinations simultaneously, but this consideration seems to be necessary in order to assign the

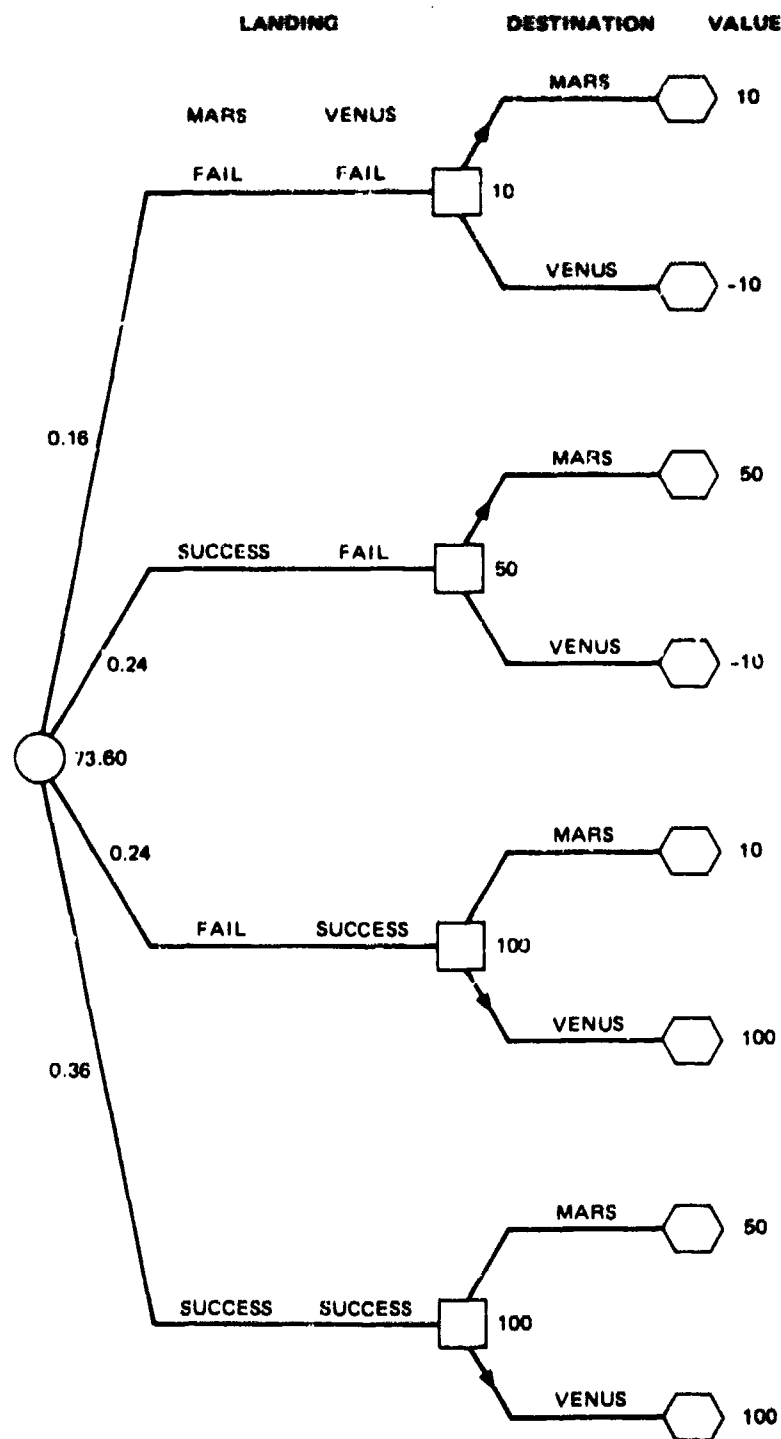


FIGURE C.3 MORE COMPLEX TREE FOR DETERMINING THE VALUE WITH PERFECT INFORMATION

required probabilities. Perhaps we could retreat to a "classical" interpretation in which we construct many "identical" hypothetical worlds where some rockets are sent to Mars and others are sent to Venus. With much thought and careful phrasing we might arrive at clear questions with useful interpretations, but we might still wonder if we had come up with a valid assessment.

### 3. The Concept of Clairvoyance

These confusing questions of probability assessment can be resolved with the introduction of the clairvoyant, a hypothetical character who knows all and who can answer any well-specified question about any uncertainty. Of course, we shall never really be able to obtain answers from him, but our probability assignments for his possible answers will provide the key to probabilistic structuring.

In the example at hand, if we were to ask the clairvoyant whether the mission will succeed or fail, he might respond by saying that the result could depend on where you sent it. Thus, we should be led to asking him two such questions, one for each destination. To make our questions precise we might draw up the questionnaire of Figure C.4. Presuming that

DESTINATION	MISSION	
	FAILURE	SUCCESS
MARS	<input type="checkbox"/>	<input type="checkbox"/>
VENUS	<input type="checkbox"/>	<input type="checkbox"/>

NOTE: Check one box in each row.

FIGURE C.4 CLAIRVOYANT'S REPORT FORM

the clairvoyant is satisfied with our definitions of success and failure, he could answer by checking one box in each row corresponding to the outcome he foretells for each destination choice.

Before we engage the clairvoyant, we wish to calculate the value of his service in monetary units: the value of clairvoyance. Since the clairvoyant has two possible answers for each of the two questions, there are four possible reports for Mars and Venus: failure, failure; failure, success; success, failure; success, success. We now must assign probabilities to these reports. A possible probability assignment, compatible with our original assignments of Figure C.1, is illustrated in Figure C.5. This distribution implies dependence between our knowledge

MARS	VENUS		
	FAILURE	SUCCESS	
FAILURE	0.354	0.046	0.4
SUCCESS	0.046	0.554	0.8
	0.4	0.6	

FIGURE C.5 JOINT DISTRIBUTION FOR CLAIRVOYANT'S ANSWERS

of the clairvoyant's two answers. For example, if he were to answer success on Mars, we would then assign a  $0.554/0.6 = 0.923$  probability that he would also answer success on Venus.

Philosophically, the important aspect of this formulation is that we are assigning probabilities to events that could occur immediately, when

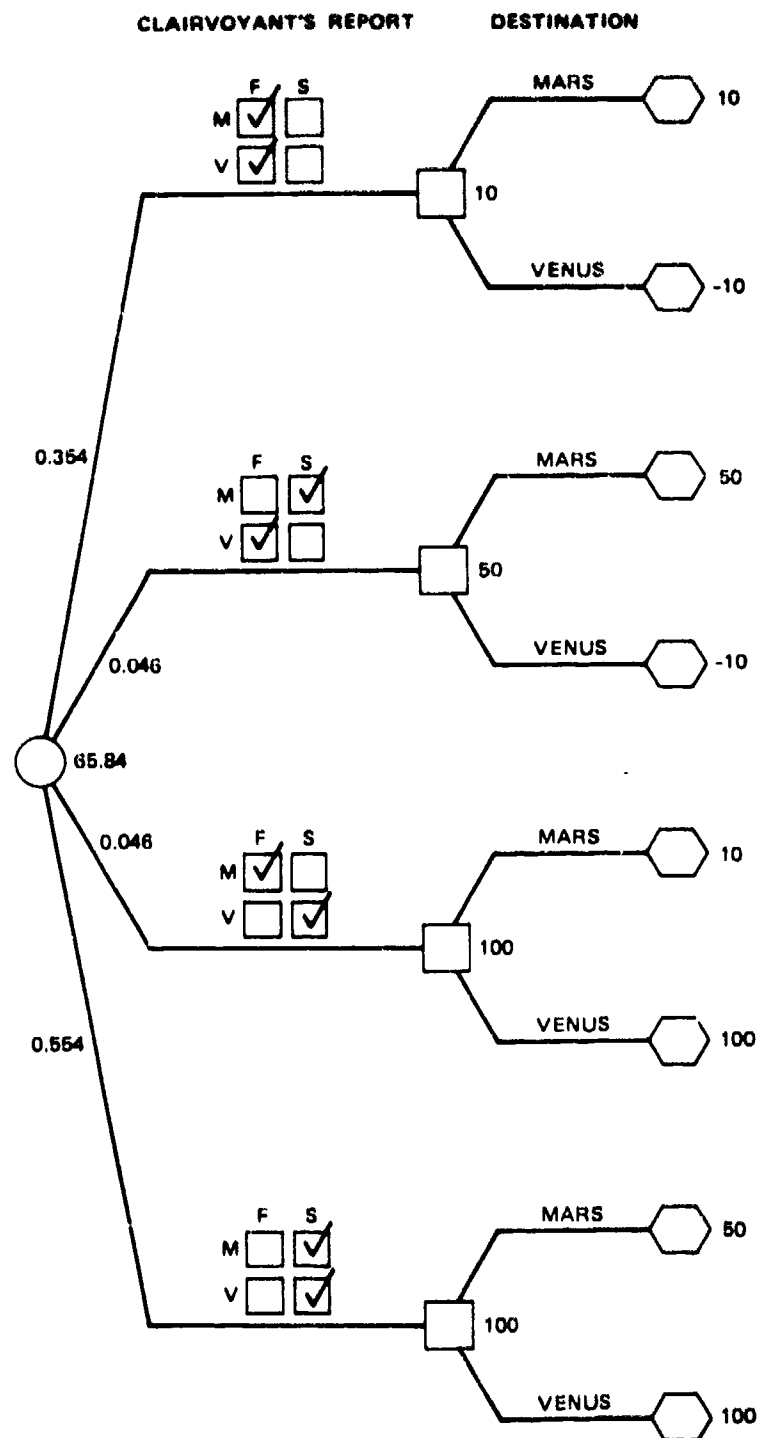


the clairvoyant reveals his answers. Also we have avoided the awkward considerations of sending our single spacecraft simultaneously to both planets or of generating hypothetical universes.

We now apply this probability assignment by constructing the decision tree of Figure C.6. The initial chance node represents the clairvoyant's revelation of one of the four possible reports, each indicated by the abbreviated report form on one of the following branches. The probabilities of Figure C.6 are assigned to these reports. Following each report we must make the best decision using the values in Figure C.1. Having made the decision indicated by the arrows, we find that the expected value with clairvoyance--but before the clairvoyant reveals his answer--is 65.84. Subtracting the 56 unit value of the primary decision (without clairvoyance) yield a value of clairvoyance of 9.84.

#### 4. Practical Probability Assignment

It would be rare for experts to think in terms of joint probability distributions such as those of Figure C.5. Normally, the experts will have technical information organized in a way that is meaningful to them, and it is desirable to construct the probability model in their terms. A simple version of such a model is illustrated by the probability assignment model of Figure C.7. Here the first question asked of the clairvoyant is, "Will the launch system, which is common to both destinations, work?" The expert has assigned 0.65 to the answer "WORK". The next two questions depend on the destination. The first is "If we launch successfully and send the spacecraft to Mars, will the landing systems work?" The expert assigns a probability of 60/65 to a positive answer. For the corresponding Venus question, the expert also assigns a probability of 60/65. (In general, these assignments need not be equal.) The expert has also stated that given a successful launch, information about the clairvoyant's report for the landing system for one destination will not influence his probability assignment for the other destination: the probability assignments to these events are conditionally independent. From this probability assignment model we can calculate the joint probability distribution for the clairvoyant's report of success or failure for the



**FIGURE C.6 DECISION TREE FOR DETERMINING THE VALUE OF COMPLETE CLAIRVOYANCE**

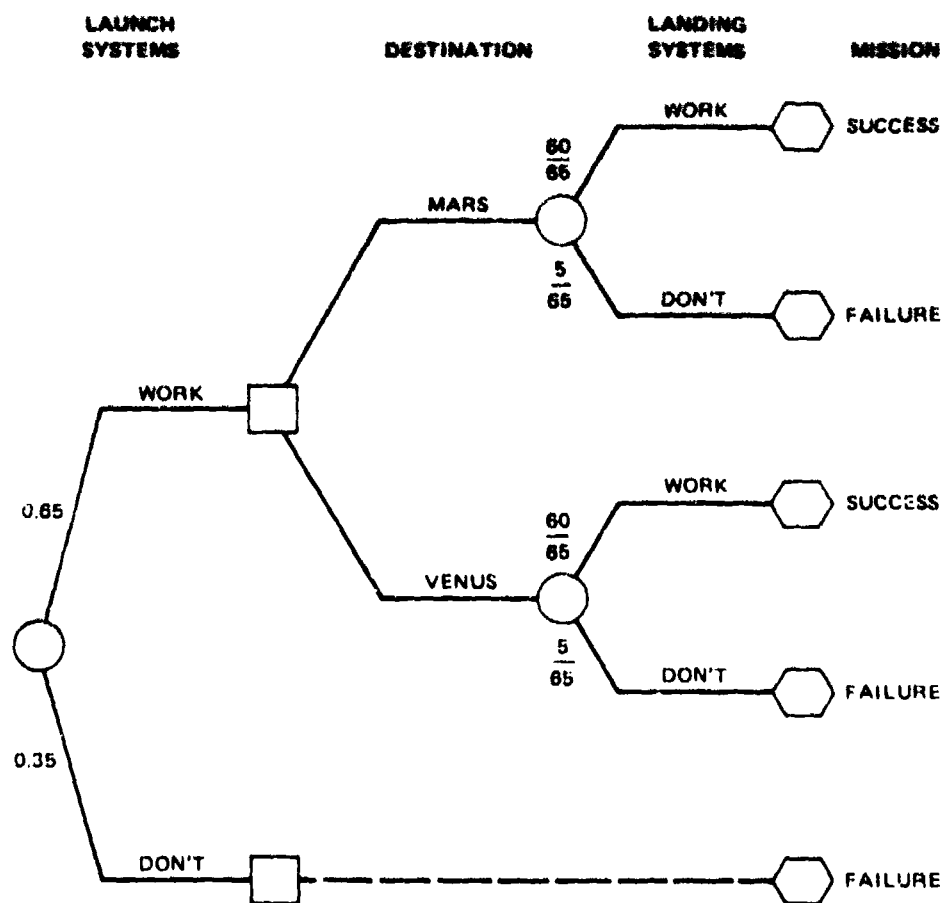


FIGURE C.7 PROBABILITY ASSIGNMENT MODEL

two possible missions (see Figure C.5). For example, the probability that the clairvoyant will report success on both planets is:

$$0.65 \frac{60}{65} \frac{60}{65} = 0.554$$

Figure C.7 is similar to what is commonly called the probability assignment tree, except that these trees usually do not include possible dependencies of probability assignments on decisions.

##### 5. Further Implications of the Probability Assignment Model

The construction of a formal probability assignment model often raises new, interesting and useful informational questions. While the

value of clairvoyance on the uncertainties appearing in the primary decision tree is an upper limit for the value of any corresponding information gathering program, often the most feasible information gathering programs are directly related to the uncertainties appearing in the probability assignment model. Thus, this model naturally leads to new and more practical information valuation questions.

For example, in the space mission problem we may be able to conduct exhaustive experiments on replicas of the launch system and make elaborate tests on the actual launch vehicle. The value of clairvoyance on the launch system alone is a straightforward calculation from the information we have built up.

The computation in Figure C.8 shows that the value with clairvoyance on the launch system is 63 units. Subtracting the 56 unit value of the primary decision, we arrive at a 7-unit value of clairvoyance on the launch system only. Since most of the 9.84 value of complete clairvoyance can be derived from the more practical launch system information, it would be best to start a realistic information-gathering program with a study of the launch system.

#### 6. Decision Dependent Clairvoyance

We may also wish to consider clairvoyance for only one of the decision alternatives. For example, suppose we engaged the clairvoyant to tell us only whether we will succeed or fail if we send the spacecraft to Mars. The primary decision tree of Figure C.1 and the probability assignment model of Figure C.7 give us all the information we need to construct the tree for the value with clairvoyance about Mars, shown in Figure C.9. Once we get the clairvoyant's report on the success of the Mars mission, we must recalculate the probability of success on Venus because of the dependency in our joint probability assignment (which results from the common launch system). A report of a successful Mars mission results in a revised probability of 0.923 for success on Venus. A report of failure if we go to Mars revises the probability of failure on Venus to 0.885. Using these probabilities, we find the expected values shown in Figure C.9. Contrary to our intuition, we find that a

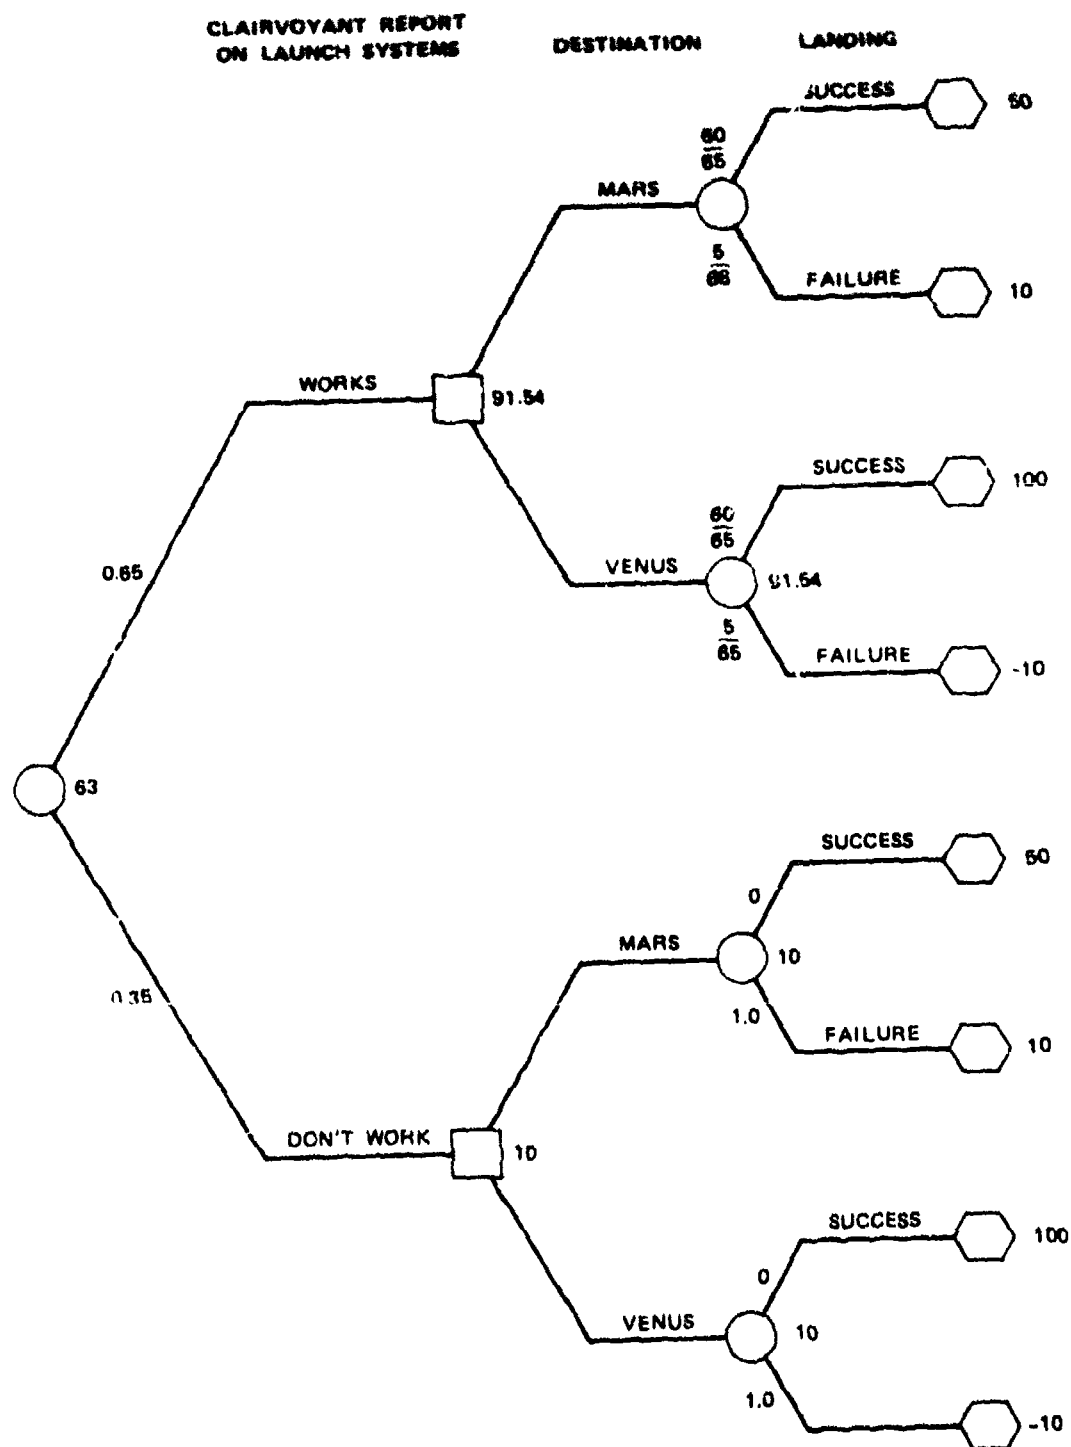
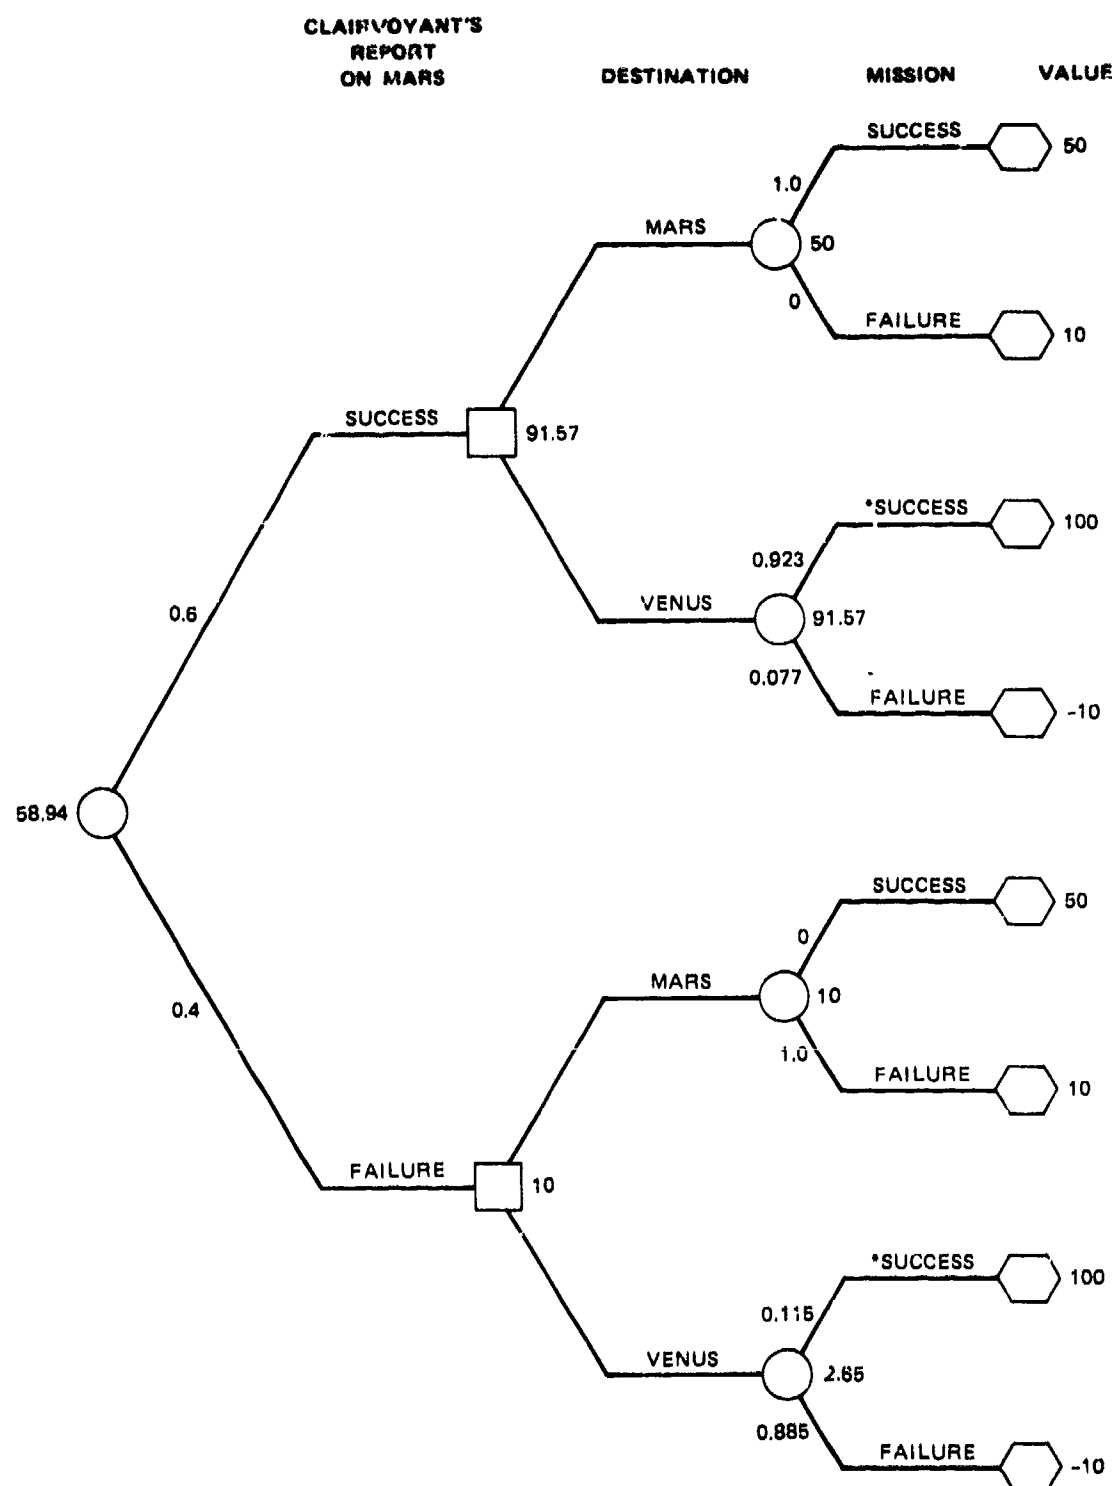


FIGURE C.8 VALUE OF CLAIRVOYANCE ON LAUNCH SYSTEMS ONLY



NOTE: \*Success probabilities on Venus are changed by Mars report.

FIGURE C.9 VALUE WITH CLAIRVOYANCE ON MARS ONLY

clairvoyant's report of a successful mars mission indicates that we should send the mission to Venus, and that a report that we will fail on Mars indicates that we should send the mission to Mars.

This is the phenomenon of decision-dependent information. Information about one aspect of a problem may have surprising implications for intuitively separate aspects of the problem due to dependencies in the probabilistic informational structure. In complex problems, a formal evaluation is the only way to determine the correct inferences and their implications.

In order to capture these effects, the analysis must not only represent the primary problem structure (Figure C.1), but it must also capture the informational structure in a formal model, (Figure C.7). In many problems the informational model may provide the more natural and more productive focus for analysis. In the spacecraft example, we can derive the primary decision tree and all the informational trees from the single probability assignment model by adding only the decision-event chronology to apply to each case. This approach to analysis might provide a key to more effective computer aids to the model building process.

#### 7. Unequal Decision-Dependent Probabilities

The calculations demonstrated above work equally well when the probabilities for success on Mars and Venus are not equal. We can demonstrate this fact by replacing the probabilities in Figure C.7 with a "launch system working" probability of 0.75, a "Mars landing system works" probability of 70/75, and a "Venus landing system works" probability of 50/75. In the primary decision tree of Figure C.1 this results in a "Mars success" probability of 0.7 and a "Venus success" probability of 0.5. The results with these new probability assignments, as well as the original ones, appear in Table C.1 along with additional clairvoyance values for the Venus mission only and for both landing systems.

#### 8. Conclusion

We have seen that a probability assessment model built on the concept of clairvoyance clarifies the interpretation and specification of

probability assignments and precisely determines values of clairvoyance. It also adds precision to the subjective interpretation and assessment

Table C.1

SUMMARY OF SPACE MISSION EXAMPLE NUMERICAL  
RESULTS WITH ORIGINAL (EQUAL) PROBABILITIES  
AND NEW (UNEQUAL) PROBABILITIES

Clairvoyance About	Equal Probabilities		Unequal Probabilities	
	Value with	Value of	Value with	Value of
Nothing	56	0	45	0
Everything	65.84	9.84	64.33	19.33
Launch system	63	7.0	50	5
Landing system	59.38	3.38	61	16
Mars only	58.94	2.94	47.34	2.34
Venus only	65.84	9.84	64.33	19.33

of probability. In problems where uncertainty plays a key role, emphasis on the construction of a formal informational model can clarify communication and lead to a more rapid and accurate solution. Further research is needed to develop a precise and convenient notational system for dealing with probability assessment models. The existing inferential notational systems are too cumbersome and the existing graphical representations are incomplete.



Appendix D

THE VALUE OF SEQUENTIAL INFORMATION

PRECEDING PAGE, BLANK, NOT FILMED

## Appendix D

### THE VALUE OF SEQUENTIAL INFORMATION

Using decision analysis it is possible to calculate the value of one of more pieces of information--called "observables"--when a decision must be made in the face of uncertainty. This information has value because it can affect the decision and lead to a better expected outcome. However, the possibility of buying information sequentially presents the decision maker with a set of secondary decisions: which observables should he buy and in which order should he buy them? It is possible that knowing one observable affects not only the primary decision, but also the decision to buy additional information. In that case the value of knowing the first observable is greater than it would be if it affected only the primary decision. The prices of the observables affect the decision maker's willingness to buy additional information. For this reason the amount that the value of learning each observable is increased by the possibility of buying additional information depends on the prices of all the observables.

When the prices of all the observables can be added to determine the price of any combination of observables and when all the prices are known with certainty, we can formulate the general sequential-information problem in terms of a set of state variables  $(x_1, \dots, x_m)$  and a set of observables  $(y_1, \dots, y_n)$  with a corresponding set of observable prices  $(K_{y_1}, \dots, K_{y_n})$ . When an observable is equal to one of the state variables it represents perfect information. However, by treating observables and state variables separately, we can also deal with imperfect information.

To solve for the value of information when all of the observables can be learned sequentially, we need to solve the decision tree shown in Figure D.1. For a large decision problem it would be very difficult and tedious to generate this decision tree. However, the tree has a very repetitive structure that can be easily implemented as part of an automated decision aid for generating decision trees. Instead of the entire

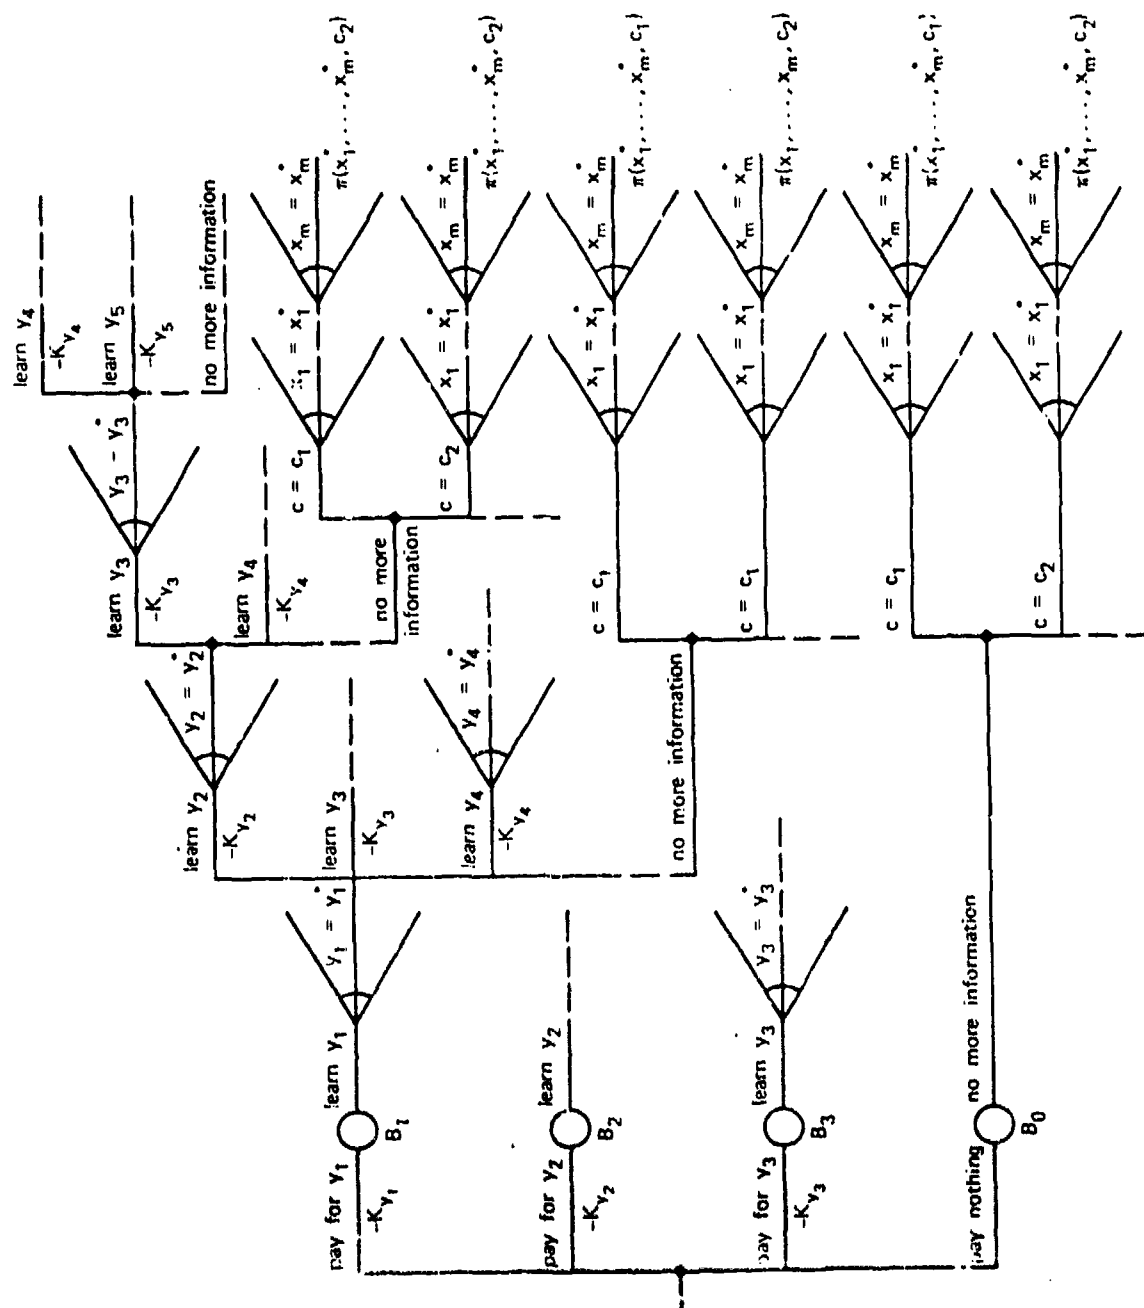


FIGURE D.1 SEQUENTIAL INFORMATION DECISION TREE

tree shown in Figure D.1, the user specifies the decision tree that exists when information cannot be purchased sequentially and then asks the computer program to expand the tree to include sequential information.

The computer program starts the expanded decision tree with a decision node such as that shown at the left of Figure D.1. The alternatives at this node are to buy any one of the specified set of observables or proceed to the basic decision tree without buying information. The last alternative leads directly to the basic decision tree specified by the user. The other alternatives lead to chance nodes where the outcome of the selected observable is revealed.

After a chance node where one of the observables is revealed, the expanded decision tree contains a decision node where the alternatives are to learn any of the specified set of observables that have not been learned previously or to buy no further information. Again, the last alternative leads directly to the basic decision tree specified by the user, except that this tree is now conditioned on the knowledge of one of the observables. The computer program continues to generate the expanded decision tree in this form until enough decision nodes are added to allow the decision maker to learn any subset of the specified observables in any order before proceeding to the primary decision problem.

By solving the decision tree in Figure D.1 several times using different prices for each of the observables, it is possible for an automated decision aid to map out decision regions such as those shown in Figure D.2. Figure D.2 shows decision regions that might occur for two observables.

We cannot regard the value of learning one observable by itself as the maximum that we would be willing to pay for that piece of information. When it is possible to buy additional information sequentially, the value of an observable may increase. To determine an upper bound for realistic efforts designed to gather sequential information, we need a decision aid that can generate and solve a decision tree like the one in Figure D.1. Without this sort of aid, the problem structuring and computations are sufficiently difficult to discourage analysts from calculating the value of sequential information, even when the results might influence information-purchasing decisions.

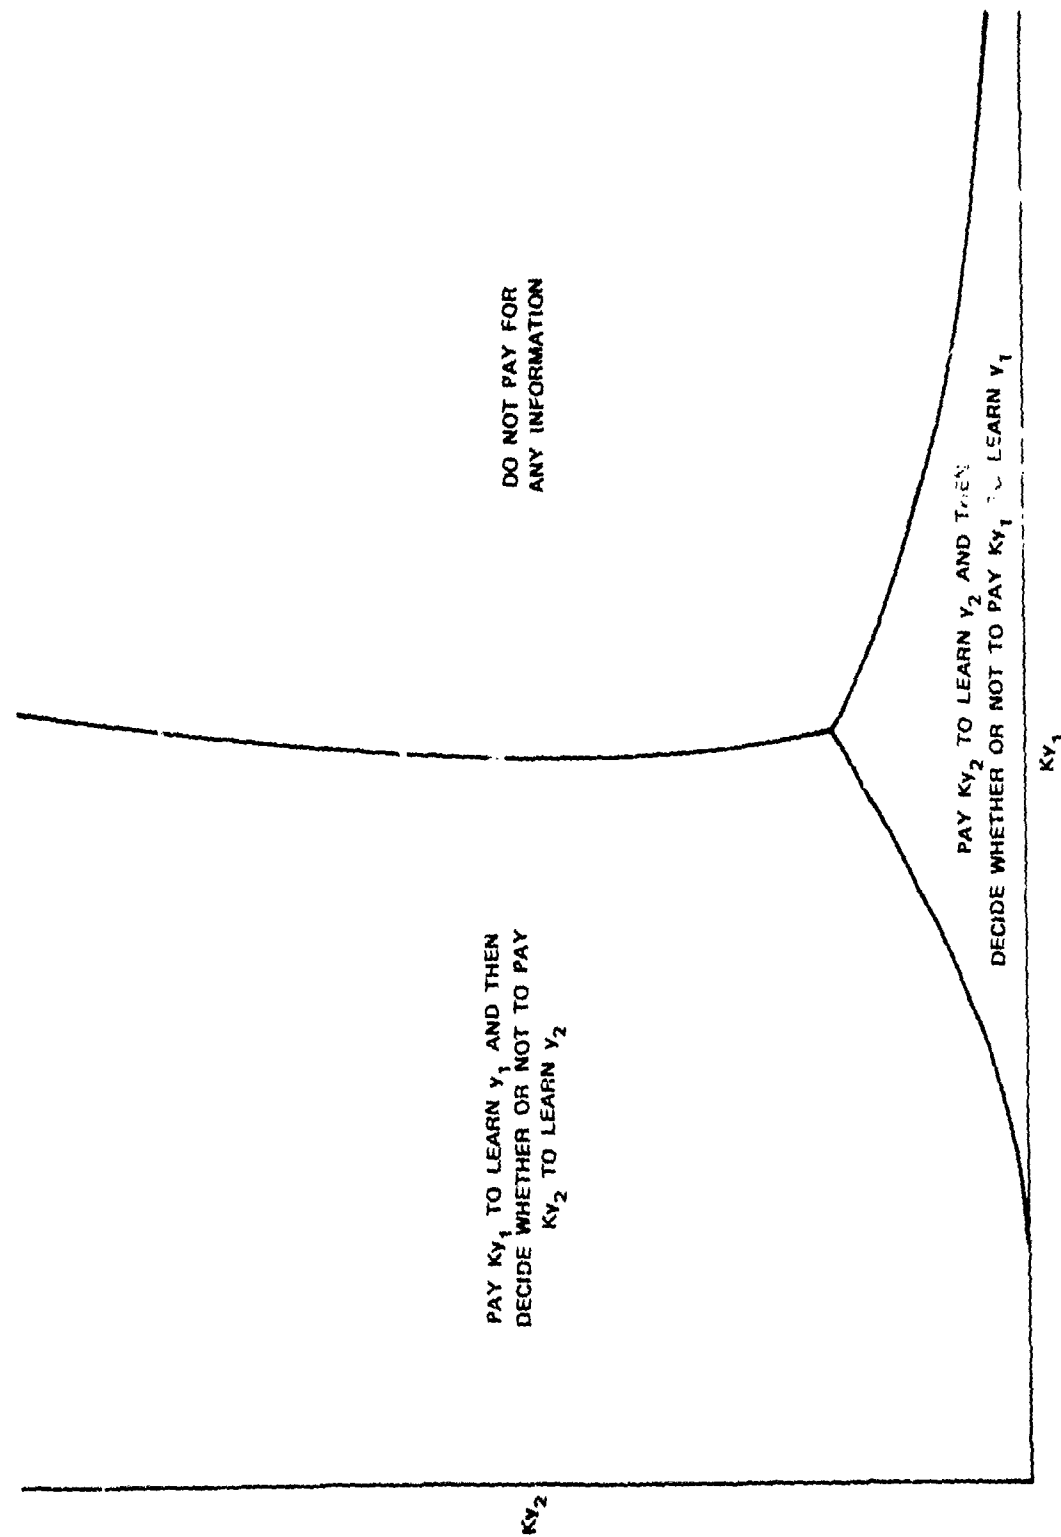


FIGURE D.2 DECISION REGIONS FOR SEQUENTIAL INFORMATION

Appendix E

THE VALUE OF FLEXIBILITY

## Appendix E

### THE VALUE OF FLEXIBILITY

The notion that a good decision strategy is a flexible one has been intuitively appreciated by decision makers for a long time. Nearly everyone is familiar with a story of a plan that went wrong because it failed to adjust for some unforeseen circumstance. Decision analysis has had little to say on the subject of flexibility. However, recent research on the concept of flexibility shows that this subject should be incorporated in a decision morphology.

Roughly speaking, something is flexible if it can be easily varied. However, in the context of decision making, ease of variation may be described by many different characteristics. Our point of view is that the flexibility of a given decision variable is determined by the nature of the choice set associated with that variable. The larger the choice set, the greater the decision flexibility. If the choice set consists of a single point element, in other words if the decision has already been committed, we say that the decision variable is inflexible.

A number of the classical micro-economic decision problems for which flexibility is a concern may be treated within this framework. Merkhofer [10] has shown that the problem of sizing a production facility can be so analyzed.\* The decision strategy of committing something less than one's total resources so as to be prepared to meet unforeseen opportunities may also be expressed as a problem of maintaining flexibility.

The value of flexibility is strongly dependent upon the information that might be received during the decision process. The more a decision maker expects to learn in the course of a decision, the more it pays to

---

\*For discussion of this problem see Marschak and Nelson [9] and Baumol [2].

follow flexible decision strategies. Similarly, the more flexible one's decision strategy, the greater the value of information gathering. Thus, the concepts of value of information and value of flexibility become special cases of the more general concept of the value of information given flexibility.

The value of information given flexibility measures the value to the decision maker, in economic units, of obtaining a given amount of information together with a given amount of decision flexibility. An upper limit to this quantity is the expected value of perfect information given perfect flexibility, EVPIGPF. Figure E.1 illustrates the calculation of the EVPIGPF in a decision tree.

Figure E.1 shows a one-stage decision problem. The decision maker must set a number of decision variables, denoted  $d_1, \dots, d_m$ . Subsequently the outcomes of a number of random variables,  $s_1, \dots, s_n$ , become known. Once he has made his decision and the information concerning the random variable is revealed, the decision maker does not have the ability to go back and alter his decision settings. The structure of the decision tree in Figure E.1a implies that our decision maker will receive no information prior to setting his decision variables.

Figure E.1b illustrates the same components of the decision problem with perfect information on state variable  $s_1$  given perfect flexibility on decision variable  $d_j$ . In this case the decision maker will learn the  $i$ 'th value of the state variable before he must set the  $j$ 'th decision variable. By calculating the maximum utility of the decision problems illustrated in both parts of Figure E.1, the value of the information given flexibility can be obtained. For an expected-value decision maker, the EVPIGPF will be the difference between the expected values associated with the two decision problems.

Thus, we see that calculation of the EVPIGPF involves the rearranging of decision and state variable nodes in the problem's decision tree. Therefore, the calculation of the value of flexibility, like the calculation of the value of information, is a tedious calculation.



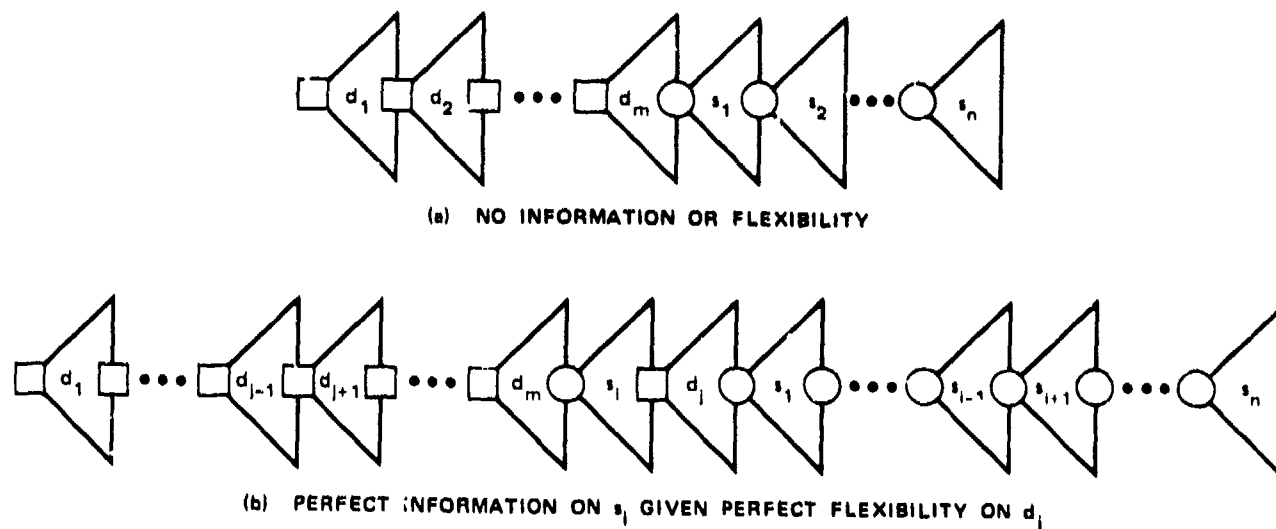


FIGURE E.1 A PORTION OF A DECISION STRUCTURE

Assistance in the form of a computerized system for restructuring decision trees would be useful for information and flexibility computations. With such an aid, the analyst would specify the potential information variables and those decision variables that could be set in response to that information. Tree restructuring would then be performed automatically and the new decision structure evaluated. The computed output would be the expected value to the decision maker of obtaining that combination of information and decision flexibility. This value would be extremely useful to the decision maker for evaluating various proposed information gathering and distribution systems.

## REFERENCES

- [1] Barrager, S.M., "Preferences for Dynamic Lotteries: Assessment and Sensitivity," Ph.D. Thesis, Stanford University, August 1975.
- [2] Baumol, W.J., "Economic Dynamics," 2nd edn., Macmillan Company, New York, 1959, pp. 92-93
- [3] Howard, R.A., "Decision Analysis: Applied Decision Theory," Proceedings of the Fourth International Conference on Operations Research (Hertz, D.B., and Melese, J., eds.), Wiley Interscience, New York, 1966.
- [4] Howard, R.A., "Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, September 1968.
- [5] Howard, R.A., "Risk Preference" (unpublished), Stanford University, 1970.
- [6] Howard, R.A., "Proximal Decision Analysis," Management Science, Vol. 17, No. 9, May 1971.
- [7] Howard, R.A., EES 221 Course Notes, Department of Engineering-Economic Systems, Stanford University, 1974.
- [8] Luce, R.D., and Raiffa, H., "Games and Decisions," Wiley, New York, 1957.
- [9] Marchak, T., and Nelson, R., "Flexibility, Uncertainty, and Economic Theory," Metroeconomica, April-December, 1962, pp. 42-58.

- [10] Merkhofer, M.W., "Flexibility and Decision Analysis," Ph.D. Dissertation, Stanford University, 1975.
- [11] Pratt, J.W., "Risk Aversion in the Small and in the Large," Econometrica, Vol. 32, January-April 1964.
- [12] Pratt, J.W., Raiffa, H., and Schlaifer, R., "Introduction to Statistical Decision Theory," McGraw-Hill, New York, 1965.
- [13] Raiffa, H., "Decision Analysis: Introductory Lectures on Choices Under Uncertainty," Addison-Wesley, Reading, Massachusetts, 1968.
- [14] Rice, T.R., "Economics of Decision Making," Research Report No. EES-DA-74-1, Department of Engineering-Economic Systems, Stanford University, August 1974.
- [15] Tribus, M., "Rational Descriptions, Decisions and Designs," Pergamon Press, New York, 1969.